arXiv:math/0603262v4 [math.AG] 22 Feb 2007

COMPUTING THE TOP BETTI NUMBERS OF SEMI-ALGEBRAIC SETS DEFINED BY QUADRATIC INEQUALITIES IN POLYNOMIAL TIME

SAUGATA BASU

Abstract. For any $\ell > 0$, we present an algorithm which takes as input a semi-algebraic set, S, defined by $P_1 \leq 0, \ldots, P_s \leq 0$, where each $P_i \in \mathrm{R}[X_1, \ldots, X_k]$ has degree ≤ 2, and computes the top ℓ Betti numbers of S, $b_{k-1}(S), \ldots, b_{k-\ell}(S)$, in polynomial time. The complexity of the algorithm, stated more precisely, is $\sum_{i=0}^{\ell+2} \binom{s}{i} k^{2^{O(\min(\ell,s))}}$. For fixed ℓ, the complexity of the algorithm can be expressed as $s^{\ell+2} k^{2^{O(\ell)}}$, which is polynomial in the input parameters s and k. To our knowledge this is the first polynomial time algorithm for computing non-trivial topological invariants of semi-algebraic sets in R^k defined by polynomial inequalities, where the number of inequalities is not fixed and the polynomials are allowed to have degree greater than one. For fixed s, we obtain by letting $\ell = k$, an algorithm for computing all the Betti numbers of S whose complexity is $k^{2^{O(s)}}$.

1. Introduction

Let R be a real closed field and $S \subset \mathrm{R}^k$ a semi-algebraic set defined by a Boolean formula with atoms of the form $P > 0, P < 0, P = 0$ for $P \in \mathcal{P} \subset \mathrm{R}[X_1, \ldots, X_k]$. It is known [23, 24, 22, 29, 4, 15] that the topological complexity of S (measured by the various Betti numbers of S) is bounded by $O(s^2 d)^k$, where $s = \#(\mathcal{P})$ and $d = \max_{P \in \mathcal{P}} \deg(P)$. Note that these bounds are singly exponential in k. More precise bounds on the individual Betti numbers of S appear in [5]. Designing efficient algorithms for computing the homology groups, and in particular the Betti numbers, of semi-algebraic sets are considered amongst the most important problems in algorithmic semi-algebraic geometry.

Even though the Betti numbers of S are bounded singly exponentially in k, there is no known algorithm for producing a singly exponential sized triangulation of S (which would immediately imply a singly exponential algorithm for computing the Betti numbers of S). In fact, the existence of a singly exponential sized triangulation, is considered to be a major open question in real algebraic geometry. Doubly exponential algorithms (with complexity $(sd)^{2^{O(k)}}$) for computing all the Betti numbers are known, since it is possible to obtain a triangulation of S in doubly exponential time using cylindrical algebraic decomposition. In the absence of singly exponential algorithms for computing triangulations of semi-algebraic sets, singly exponential algorithms are known only for the problems of testing emptiness

Key words and phrases. Betti numbers, Quadratic inequalities, Semi-algebraic sets.

The author was supported in part by an NSF Career Award 0133597 and a Sloan Foundation Fellowship. A preliminary version of this paper appears in the Proceedings of the ACM Symposium on Theory of Computing, 2005.

[27, 7], computing the zero-th Betti number (i.e. the number of semi-algebraically connected components of S) [18, 20, 14, 16, 8], as well as the Euler-Poincaré characteristic of S [4]. Very recently a singly exponential time algorithm has been given for computing the first few Betti numbers of semi-algebraic sets [3] (see also [6]).

In this paper, we consider a restricted class of semi-algebraic sets – namely, semi-algebraic sets defined by a conjunction of quadratic inequalities. Since sets defined by linear inequalities have no interesting topology, sets defined by quadratic inequalities can be considered to be the simplest class of semi-algebraic sets which can have non-trivial topology. Such sets are in fact quite general, since every semi-algebraic set can be defined by (quantified) formulas involving only quadratic polynomials (at the cost of increasing the number of variables and the size of the formula). Moreover, as in the case of general semi-algebraic sets, the Betti numbers of such sets can be exponentially large. For example, the set $S \subset \mathrm{R}^k$ defined by

$$X_1(1-X_1) \leq 0, \ldots, X_k(1-X_k) \leq 0,$$

has $b_0(S) = 2^k$.

Hence, it is somewhat surprising that for any fixed constant ℓ, the Betti numbers $b_{k-1}(S), \ldots, b_{k-\ell}(S)$, of a basic closed semi-algebraic set $S \subset \mathrm{R}^k$ defined by quadratic inequalities, are polynomially bounded. The following theorem appears in [5].

Theorem 1.1. *Let* R *a real closed field and* $S \subset \mathrm{R}^k$ *be defined by*

$$P_1 \leq 0, \ldots, P_s \leq 0, \deg(P_i) \leq 2, 1 \leq i \leq s.$$

Then, for $\ell \geq 0$*,*

$$b_{k-\ell}(S) \leq \binom{s}{\ell} k^{O(\ell)}.$$

Notice that for fixed ℓ this gives a polynomial bound on the highest ℓ Betti numbers of S (which could possibly be non-zero). Observe also that similar bounds do not hold for sets defined by polynomials of degree greater than two. For instance, the set $V \subset \mathrm{R}^k$ defined by the single quartic inequality,

$$\sum_{i=1}^{k} X_i^2(X_i-1)^2 - \varepsilon \geq 0,$$

will have $b_{k-1}(V) = 2^k$, for all small enough $\varepsilon > 0$.

To see this observe that for all sufficiently small $\varepsilon > 0$, $\mathrm{R}^k \setminus V$ is defined by

$$\sum_{i=1}^{k} X_i^2(X_i-1)^2 < \varepsilon.$$

and has 2^k connected components, since it retracts onto the set $\{0,1\}^k$. It now follows by Alexander duality that

$$b_{k-1}(V) = b_0(\mathrm{R}^k \setminus V) = 2^k.$$

Remark 1.2. Even though Theorem 1.1 is stated for semi-algebraic sets defined by a conjunction of weak inequalities, there is an easy reduction to this case for basic semi-algebraic sets defined by equalities and strict inequalities. The same reduction is also applicable to the algorithmic results described later in this paper, in particular Theorem 6.2.

Note that the definition of cohomology groups of basic semi-algebraic sets defined by equalities and strict inequalities over an arbitrary real closed field R requires some care, and several possibilities exist. In this paper we follow the definition given in [9] which agrees with singular cohomology in case $\mathrm{R} = \mathbb{R}$. We refer the reader to [9] for further elaboration on this point.

We have the following easy corollary of Theorem 1.1.

Corollary 1.3. *Let* R *a real closed field and* $S \subset \mathrm{R}^k$ *be defined by*

$$\bigwedge_{P \in \mathcal{P}_1} P = 0 \bigwedge_{P \in \mathcal{P}_2} P > 0 \bigwedge_{P \in \mathcal{P}_3} P < 0$$

with $\deg(P) \leq 2$ *for each* $P \in \bigcup_{i=0,1,2} \mathcal{P}_i$*, and* $\# \bigcup_{i=0,1,2} \mathcal{P}_i = s$*.*

Then, for all $\ell \geq 0$*,*

$$b_{k-\ell}(S) \leq \binom{s}{\ell} k^{O(\ell)}.$$

In the following proof, as well as later in the paper, we will extend the ground field R by infinitesimal elements. We denote by $\mathrm{R}\langle\zeta\rangle$ the real closed field of algebraic Puiseux series in ζ with coefficients in R (see [9] for more details). The sign of a Puiseux series in $\mathrm{R}\langle\zeta\rangle$ agrees with the sign of the coefficient of the lowest degree term in ζ. This induces a unique order on $\mathrm{R}\langle\zeta\rangle$ which makes ζ infinitesimal: ζ is positive and smaller than any positive element of R. When $a \in \mathrm{R}\langle\zeta\rangle$ is bounded from above and below by some elements of R, $\lim_\zeta(a)$ is the constant term of a, obtained by substituting 0 for ζ in a. Given a semi-algebraic set S in R^k, the *extension* of S to R', denoted $\mathrm{Ext}(S, \mathrm{R}')$, is the semi-algebraic subset of ${\mathrm{R}'}^k$ defined by the same quantifier free formula that defines S. The set $\mathrm{Ext}(S, \mathrm{R}')$ is well defined (i.e. it only depends on the set S and not on the quantifier free formula chosen to describe it). This is an easy consequence of the transfer principle (see for instance [9]).

Proof of Corollary 1.3. Let $0 < \delta \ll \varepsilon \ll 1$ be infinitesimals. We first replace the set S by the set $S' \subset \mathrm{R}\langle\varepsilon\rangle^k$ defined by $S' = \mathrm{Ext}(S, \mathrm{R}\langle\varepsilon\rangle) \cap \bar{B}_k(0, 1/\varepsilon)$, where $\bar{B}_k(0, r)$ denotes the closed ball of radius r centered at the origin. It follows from Hardt's triviality theorem for semi-algebraic mappings [19] that $b_i(S) = b_i(S')$ for all $i \geq 0$. We then replace S' by the set $S'' \subset \mathrm{R}\langle\varepsilon, \delta\rangle^k$ defined by,

$$\bigwedge_{P \in \mathcal{P}_1} P \leq 0 \wedge -P \leq 0 \bigwedge_{P \in \mathcal{P}_2} -P + \delta \leq 0 \bigwedge_{P \in \mathcal{P}_3} -P - \delta \leq 0 \bigwedge \varepsilon^2(X_1^2 + \cdots + X_k^2) - 1 \leq 0.$$

It follows from Hardt's triviality again that, $b_i(S') = b_i(S'')$ for all $i \geq 0$. Now apply Theorem 1.1. □

Remark 1.4. Another point to note is that the bound in Theorem 1.1 depends crucially on the structure of the formula defining the semi-algebraic set S – namely that S is defined by a *conjunction* of polynomial inequalities of the form $P \leq 0$ with $\deg(P) \leq 2$. The polynomial bound on the highest Betti numbers no longer holds if S is defined by a formula involving quadratic inequalities but having a different structure – for instance, a disjunction instead of a conjunction of such inequalities. However, the polynomial bound continues to hold for certain small variations of the structure of the formula defining S. For instance, we have the following slight generalization of Theorem 1.1 whose proof mimics that of Theorem 1.1 and which we omit.

Theorem 1.5. *Let* R *a real closed field and* $S \subset \mathrm{R}^k$ *be defined by*

$$\bigwedge_{1\leq i\leq s} (\bigvee_{1\leq j\leq m} P_{ij} \leq 0), \deg(P_{ij}) \leq 2, 1 \leq i \leq s, 1 \leq j \leq m.$$

Then, for $\ell \geq 0$*,*

$$b_{k-\ell}(S) \leq \binom{s}{\ell} k^{O(m\ell)}.$$

For fixed m and ℓ, the above bound is still polynomial in s and k. Similarly, the main result of this paper, namely Algorithm 4 in Section 8, can be extended to this situation as well with complexity polynomial in s and k for fixed m and ℓ. We will omit the details of this extension.

Semi-algebraic sets defined by a system of quadratic inequalities have a special significance in the theory of computational complexity. Even though such sets might seem to be the next simplest class of semi-algebraic sets after sets defined by linear inequalities, from the point of view of computational complexity they represent a quantum leap. Whereas there exist (weakly) polynomial time algorithms for solving linear programming, solving quadratic feasibility problem is provably hard. For instance, it follows from an easy reduction from the problem of testing feasibility of a real quartic equation in many variables, that the problem of testing whether a system of quadratic inequalities is feasible is NP_{R}-complete in the Blum-Shub-Smale model of computation (see [10]). Assuming the input polynomials to have integer coefficients, the same problem is NP-hard in the classical Turing machine model, since it is also not difficult to see that the Boolean satisfiability problem can be posed as the problem of deciding whether a certain semi-algebraic set defined by quadratic inequalities is empty or not (see Section 9). Counting the number of connected components of such sets is even harder. In fact, we prove (see Theorem 9.1) that for $\ell = O(\log k)$, computing the ℓ-th Betti number of a basic semi-algebraic set defined by quadratic inequalities in R^k is #P-hard. Note that PSPACE-hardness of the problem of counting the number of connected components for general semi-algebraic sets were known before [13, 26], and the proofs of these results extend easily to the quadratic case. In view of these hardness results, it is unlikely that there exist polynomial time algorithms for computing the Betti numbers (or even the first few Betti numbers) of such a set. In contrast to these hardness results, the polynomial bound on the top Betti numbers of sets defined by quadratic inequalities gives rise to the possibility that these might in fact be computable in polynomial time.

In this paper we prove that for each fixed $\ell > 0$, the top ℓ Betti numbers of basic semi-algebraic sets defined by quadratic inequalities are computable in polynomial time. We will assume that the polynomials given as input to our algorithms have coefficients in some ordered domain D contained in a real closed field R. We will denote the algebraic closure of R by C. By complexity of our algorithms we will mean the number of arithmetic operations including comparisons in the ring D. When $\mathrm{D} = \mathbb{Z}$, we will also count the number of bit operations.

The main result of this paper is the following.

Main Result: We present an algorithm (Algorithm 4 in Section 8) which given a set of s polynomials, $\mathcal{P} = \{P_1, \ldots, P_s\} \subset \mathrm{R}[X_1, \ldots, X_k]$, with $\deg(P_i) \leq 2, 1 \leq i \leq s$, computes $b_{k-1}(S), \ldots, b_{k-\ell}(S)$, where S is the set defined by $P_1 \leq 0, \ldots, P_s \leq 0$.

The complexity of the algorithm is

$$\sum_{i=0}^{\ell+2} \binom{s}{i} k^{2^{O(\min(\ell,s))}}. \tag{1.1}$$

If the coefficients of the polynomials in $\mathcal{P}$ are integers of bitsizes bounded by τ, then the bitsizes of the integers appearing in the intermediate computations and the output are bounded by $\tau(sk)^{2^{O(\min(\ell,s))}}$.

To our knowledge this is the first polynomial time algorithm for computing a non-trivial topological invariant of semi-algebraic sets in R^k defined by polynomial inequalities, where the number of constraints is not fixed and the polynomials are allowed to have degree greater than 1. The special case when all the polynomials are linear reduces to the well-studied problem of linear programming. In this case the set S is either a convex polyhedron or empty, and (weakly) polynomial time algorithms are known to decide emptiness of such a set. In another direction, Barvinok [2] designed a polynomial time algorithm for deciding feasibility of systems of quadratic inequalities, but under the condition that the number of inequalities is bounded by a constant (see also [17] for an interesting generalization as well as a constructive version of this result).

2. Brief Outline

Given any compact semi-algebraic set S, we will denote by $b_i(S)$ the rank of $H^i(S,\mathbb{Q})$ (the i-th simplicial cohomology group of S with coefficients in $\mathbb{Q}$). We denote by $\mathbf{S}^k \subset \mathrm{R}^{k+1}$ the unit sphere centered at the origin. We first consider the case of semi-algebraic subsets of the unit sphere, $\mathbf{S}^k \subset \mathrm{R}^{k+1}$, defined by homogeneous quadratic inequalities. We then show how to reduce the general problem to this special case.

Let $S \subset \mathbf{S}^k$ be the set defined on $\mathbf{S}^k$ by s inequalities, $P_1 \leq 0, \ldots, P_s \leq 0$, where $P_1, \ldots, P_s \in \mathrm{R}[X_0, \ldots, X_k]$ are homogeneous quadratic polynomials. For each $i, 1 \leq i \leq s$, let $S_i \subset \mathbf{S}^k$ denote the set defined on $\mathbf{S}^k$ by $P_i \leq 0$. Then, $S = \cap_{i=1}^s S_i$. There are two main ingredients in the polynomial time algorithm for computing the top Betti numbers of S.

The first main idea is to consider S as the intersection of the various S_i's and to utilize the double complex arising from the generalized Mayer-Vietoris exact sequence (see Section 3). It follows from the exactness of the generalized Mayer-Vietoris sequence (see Proposition 1 below), that the top dimensional homology groups of S are isomorphic to those of the total complex associated to a suitable truncation of the Mayer-Vietoris double complex. However, computing even the truncation of the Mayer-Vietoris double complex, starting from a triangulation of S would entail a doubly exponential complexity. However, we utilize the fact that terms appearing in the truncated complex depend on the unions of the S_i's taken at most $\ell + 2$ at a time. There are at most $\sum_{j=1}^{\ell+2} \binom{s}{j}$ such sets. Moreover, for semi-algebraic sets defined by the disjunction of a small number of quadratic inequalities, we are able to compute in polynomial (in k) time a complex, whose homology groups are isomorphic to those of the given sets. The construction of these complexes in polynomial time is the second important ingredient in our algorithm and is described in detail in Section 7 (Algorithm 2). These complexes along with the homomorphisms between them define another double complex whose associated spectral sequence (corresponding to the column-wise filtration) is isomorphic from

the E_2 term onwards to the corresponding one of the (truncated) Mayer-Vietoris double complex (see Theorem 6.2 below). Since, we know that the latter converges to the homology groups of S, the top Betti numbers of S are equal to the ranks of the homology groups of the associated total complex of the double complex we computed. These can then be computed using well known efficient algorithms from linear algebra.

The rest of the paper is organized as follows. In Section 3 we recall certain basic facts from algebraic topology including the notions of complexes, and double complexes of vector spaces, spectral sequences and triangulations of semi-algebraic sets. We do not prove any results since all of them are quite classical and we refer the reader to appropriate references [12, 21, 9] for the proofs. In Section 4, we recall some basic algorithms in semi-algebraic geometry that we will need later. We state the inputs, outputs and complexities of these algorithms. For classical results in the field of real algebraic geometry, as well as the details of certain algorithms that we use, we refer the reader to [9] for convenience. In Section 5 we describe certain topological properties of semi-algebraic sets defined by quadratic inequalities which are crucial for our algorithm. Most of the results in this section are due to Agrachev [1]. In Section 6 we prove the main mathematical results necessary for our algorithm. In Section 7 we describe our algorithm for computing the top Betti numbers of semi-algebraic sets defined by quadratic forms. We treat the general case in Section 8. Finally, in Section 9, we show the computational hardness of computing the first few Betti numbers of a given semi-algebraic set defined by quadratic inequalities, by proving that the problem is #P-hard.

3. Topological Preliminaries

In this section we recall some basic facts from algebraic topology, related to double complexes, and spectral sequences associated to double complexes. We also fix our notations for these objects. All facts that we need are well known, and we merely give a brief overview, referring the reader to [12, 21] for detailed proofs.

3.1. **Complex of Vector Spaces.** A sequence $\{C^p\}$, $p \in \mathbb{Z}$, of $\mathbb{Q}$-vector spaces together with a sequence $\{\partial^p\}$ of homomorphisms $\partial^p : C^p \to C^{p+1}$ for which $\partial^{p+1} \circ \partial^p = 0$ for all p is called a complex.

The cohomology groups, $H^p(C^\bullet, \mathbb{Q})$ are defined by,

$$H^p(C^\bullet, \mathbb{Q}) = Z^p(C^\bullet)/B^p(C^\bullet),$$

where $B^p(C^\bullet) = \mathrm{Im}(\partial^{p-1})$, and $Z^p(C^\bullet) = \mathrm{Ker}(\partial^p)$.

The cohomology groups, $H^*(C^\bullet, F)$, are all $\mathbb{Q}$-vector spaces (finite dimensional if the vector spaces C^p's are themselves finite dimensional). We will henceforth omit reference to the field of coefficients $\mathbb{Q}$ which is fixed throughout the rest of the paper.

Given a complex $C^\bullet$, we denote by $\check{C}_\bullet$ the dual complex,

$$\cdots \longleftarrow \check{C}_{p-1} \xleftarrow{\check{\partial}_p} \check{C}_p \longleftarrow \cdots$$

where $\check{C}_p = \mathrm{Hom}(C^p, \mathbb{Q})$ is the vector space dual to C^p and

$$\check{\partial}_p : \mathrm{Hom}(C^p, \mathbb{Q}) \to \mathrm{Hom}(C^{p-1}, \mathbb{Q})$$

is the homomorphism dual to ∂^{p-1}.

Moreover, $H_*(\check{C}_\bullet, \mathbb{Q}) \cong H^*(C^\bullet, \mathbb{Q})$.

Given two complexes, $C^\bullet = (C^p, \partial^p)$ and $D^\bullet = (D^p, \partial^p)$, a homomorphism of complexes, $\phi : C^\bullet \to D^\bullet$, is a sequence of homomorphisms $\phi^p : C^p \to D^p$ for which $\partial^p \circ \phi^p = \phi^{p+1} \circ \partial^p$ for all p.

In other words, the following diagram is commutative.

$$\begin{array}{ccccccc} \cdots & \longrightarrow & C^p & \xrightarrow{\partial^p} & C^{p+1} & \longrightarrow & \cdots \\ & & \downarrow \phi^p & & \downarrow \phi^{p+1} & & \\ \cdots & \longrightarrow & D^p & \xrightarrow{\partial^p} & D^{p+1} & \longrightarrow & \cdots \end{array}$$

A homomorphism of complexes, $\phi : C^\bullet \to D^\bullet$, induces homomorphisms, $\phi^* : H^*(C^\bullet) \to H^*(D^\bullet)$. The homomorphism ϕ is called a *quasi-isomorphism* if the homomorphisms ϕ^* are isomorphisms.

3.2. **Double Complexes.** A *double complex* is a bi-graded vector space,

$$C^{\bullet,\bullet} = \bigoplus_{p,q \in \mathbb{Z}} C^{p,q},$$

with co-boundary operators $d : C^{p,q} \to C^{p,q+1}$ and $\delta : C^{p,q} \to C^{p+1,q}$ and such that $d\delta + \delta d = 0$. We say that $C^{\bullet,\bullet}$ is a first quadrant double complex, if it satisfies the condition that $C^{p,q} = 0$ if either $p < 0$ or $q < 0$. Double complexes lying in other quadrants are defined in an analogous manner.

$$\begin{array}{ccccccc} \vdots & & \vdots & & \vdots & & \\ \uparrow d & & \uparrow d & & \uparrow d & & \\ C^{0,2} & \xrightarrow{\delta} & C^{1,2} & \xrightarrow{\delta} & C^{2,2} & \xrightarrow{\delta} & \cdots \\ \uparrow d & & \uparrow d & & \uparrow d & & \\ C^{0,1} & \xrightarrow{\delta} & C^{1,1} & \xrightarrow{\delta} & C^{2,1} & \xrightarrow{\delta} & \cdots \\ \uparrow d & & \uparrow d & & \uparrow d & & \\ C^{0,0} & \xrightarrow{\delta} & C^{1,0} & \xrightarrow{\delta} & C^{2,0} & \xrightarrow{\delta} & \cdots \end{array}$$

We call the complex, $\mathrm{Tot}^\bullet(C^{\bullet,\bullet})$, defined by

$$\mathrm{Tot}^n(C^{\bullet,\bullet}) = \bigoplus_{p+q=n} C^{p,q},$$

with differential

$$\mathrm{D}^n = d + \delta : \mathrm{Tot}^n(C^{\bullet,\bullet}) \to \mathrm{Tot}^{n+1}(C^{\bullet,\bullet}),$$

to be the *associated total complex of* $C^{\bullet,\bullet}$.

$$\begin{tikzcd}
 & \vdots & \vdots & \vdots & \\
\cdots \arrow[r, "\delta"] & C^{p-1,q+1} \arrow[r, "\delta"] \arrow[u, "d"] & C^{p,q+1} \arrow[r, "\delta"] \arrow[u, "d"] & C^{p+1,q+1} \arrow[r, "\delta"] \arrow[u, "d"] & \cdots \\
\cdots \arrow[r, "\delta"] & C^{p-1,q} \arrow[r, "\delta"] \arrow[u, "d"] & C^{p,q} \arrow[r, "\delta"] \arrow[u, "d"] & C^{p+1,q} \arrow[r, "\delta"] \arrow[u, "d"] & \cdots \\
\cdots \arrow[r, "\delta"] & C^{p-1,q-1} \arrow[r, "\delta"] \arrow[u, "d"] & C^{p,q-1} \arrow[r, "\delta"] \arrow[u, "d"] & C^{p+1,q-1} \arrow[r, "\delta"] \arrow[u, "d"] & \cdots \\
 & \vdots \arrow[u, "d"] & \vdots \arrow[u, "d"] & \vdots \arrow[u, "d"] &
\end{tikzcd}$$

3.3. **Spectral Sequences.** A *spectral sequence* is a sequence of bigraded complexes $(E_r, d_r : E_r^{p,q} \to E_r^{p+r,q-r+1})$ (see Figure 1) such that the complex E_{r+1} is obtained from E_r by taking its homology with respect to d_r (that is $E_{r+1} = H_{d_r}(E_r)$).

There are two spectral sequences, $'E_*^{p,q}, ''E_*^{p,q}$, (corresponding to taking row-wise or column-wise filtrations respectively) associated with a double complex $C^{\bullet,\bullet}$, which will be important for us. Both of these converge to $H^*(\mathrm{Tot}^\bullet(C^{\bullet,\bullet}))$. This means that the homomorphisms d_r are eventually zero, and hence the spectral sequences stabilize, and

$$\bigoplus_{p+q=i} {}'E_\infty^{p,q} \cong \bigoplus_{p+q=i} {}''E_\infty^{p,q} \cong H^i(\mathrm{Tot}^\bullet(C^{\bullet,\bullet})),$$

for each $i \geq 0$.

The first terms of these are $'E_1 = H_\delta(C^{\bullet,\bullet}), 'E_2 = H_d H_\delta(C^{\bullet,\bullet})$, and $''E_1 = H_d(C^{\bullet,\bullet}), ''E_2 = H_\delta H_d(C^{\bullet,\bullet})$.

Given two (first quadrant) double complexes, $C^{\bullet,\bullet}$ and $\bar{C}^{\bullet,\bullet}$, a homomorphism of double complexes,

$$\phi : C^{\bullet,\bullet} \longrightarrow \bar{C}^{\bullet,\bullet},$$

is a collection of homomorphisms, $\phi^{p,q} : C^{p,q} \longrightarrow \bar{C}^{p,q}$, such that the following diagrams commute.

$$\begin{tikzcd}
C^{p,q} \arrow[r, "\delta"] \arrow[d, "\phi^{p,q}"] & C^{p+1,q} \arrow[d, "\phi^{p+1,q}"] \\
\bar{C}^{p,q} \arrow[r, "\delta"] & \bar{C}^{p+1,q}
\end{tikzcd}$$

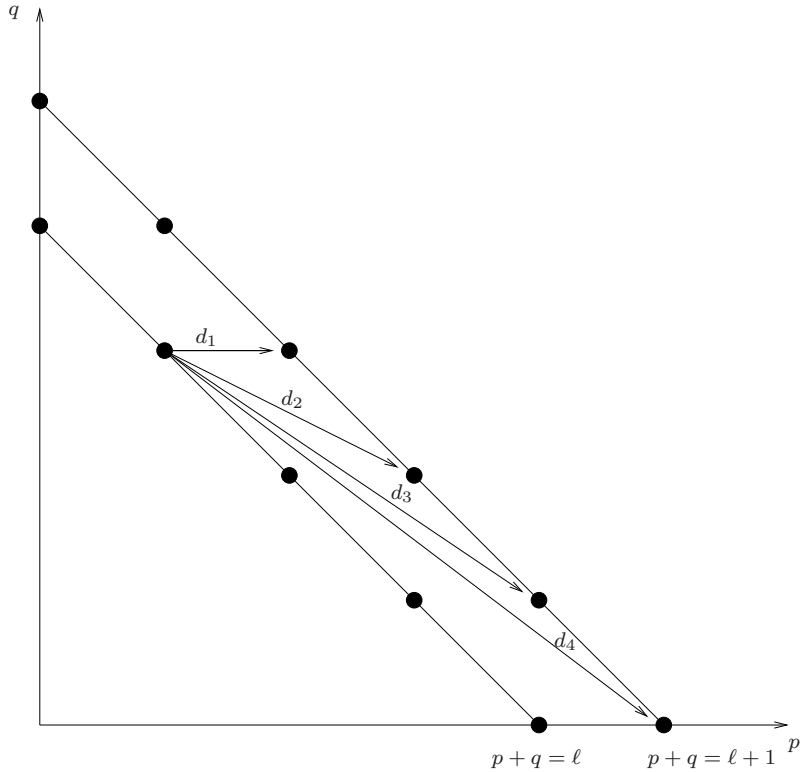

FIGURE 1. $d_r : E_r^{p,q} \to E_r^{p+r,q-r+1}$

$$\begin{array}{ccc} C^{p,q} & \xrightarrow{d} & C^{p,q+1} \\ \downarrow \phi^{p,q} & & \downarrow \phi^{p,q+1} \\ \bar{C}^{p,q} & \xrightarrow{d} & \bar{C}^{p,q+1} \end{array}$$

A homomorphism of double complexes,

$$\phi : C^{\bullet,\bullet} \longrightarrow \bar{C}^{\bullet,\bullet},$$

induces homomorphisms, $\phi_s : E_s \longrightarrow \bar{E}_s$ between the associated spectral sequences (corresponding either to the row-wise or column-wise filtrations). For the precise definition of homomorphisms of spectral sequences, see [21]. We will use the following useful fact (see [21], page 66, Theorem 3.4) several times in the rest of the paper.

Theorem 3.1. *If ϕ_s is an isomorphism for some $s \geq 1$, then $E_r^{p,q}$ and $\bar{E}_r^{p,q}$ are isomorphic for all $r \geq s$. In other words, the induced homomorphism, ϕ : $\mathrm{Tot}^\bullet(C^{\bullet,\bullet}) \to \mathrm{Tot}^\bullet(\bar{C}^{\bullet,\bullet})$ is a quasi-isomorphism.*

3.4. **Triangulation of semi-algebraic sets.** A triangulation of a compact semi-algebraic set S is a simplicial complex Δ together with a semi-algebraic homeomorphism from $|\Delta|$ to S. Given such a triangulation we will often identify the simplices

in Δ with their images in S under the given homeomorphism, and will refer to the triangulation by Δ.

Given a triangulation Δ, the cohomology groups $H^i(S)$ are isomorphic to the simplicial cohomology groups $H^i(\Delta)$ of the simplicial complex Δ and are in fact independent of the triangulation Δ (this fact is classical over $\mathbb{R}$; see for instance [9] for a self-contained proof in the category of semi-algebraic sets).

We call a triangulation $h_1 : |\Delta_1| \to S$ of a semi-algebraic set S, to be a *refinement* of a triangulation $h_2 : |\Delta_2| \to S$ if for every simplex $\sigma_1 \in \Delta_1$, there exists a simplex $\sigma_2 \in \Delta_2$ such that $h_1(\sigma_1) \subset h_2(\sigma_2)$.

Let $S_1 \subset S_2$ be two compact semi-algebraic subsets of R^k. We say that a semi-algebraic triangulation $h : |\Delta| \to S_2$ of S_2, *respects* S_1 if for every simplex $\sigma \in \Delta$, $h(\sigma) \cap S_1 = h(\sigma)$ or $\emptyset$. In this case, $h^{-1}(S_1)$ is identified with a sub-complex of Δ and $h|_{h^{-1}(S_1)} : h^{-1}(S_1) \to S_1$ is a semi-algebraic triangulation of S_1. We will refer to this sub-complex by $\Delta|_{S_1}$.

We will need the following theorem which can be deduced from Section 9.2 in [11] (see also [9]).

Theorem 3.2. *Let $S_1 \subset S_2 \subset \mathrm{R}^k$ be closed and bounded semi-algebraic sets, and let $h_i : \Delta_i \to S_i, i = 1, 2$ be semi-algebraic triangulations of S_1, S_2. Then, there exists a semi-algebraic triangulation $h : \Delta \to S_2$ of S_2, such that Δ respects S_1, Δ is a refinement of Δ_2, and $\Delta|_{S_1}$ is a refinement of Δ_1.*

Moreover, there exists an algorithm which computes such a triangulation whose complexity is bounded by $(sd)^{2^{O(k)}}$, where s is the number of polynomials used in the definition of S_1 and S_2, and d a bound on their degrees.

3.5. **Mayer-Vietoris Double Complex.** Let $S_1, \ldots, S_s \subset \mathrm{R}^k$ be closed and bounded semi-algebraic sets, and let $S = \cap_{i=1}^s S_i$. Choose a sufficiently fine triangulation of $\cup_{1=1}^s S_i$, such that all intersections of the form, $S_{i_0} \cap \cdots \cap S_{i_\ell}$, correspond to subcomplexes of the simplicial complex of the triangulation. Note that the existence of such a triangulation (in fact, a semi-algebraic triangulation) is well known (see [9]). However, the best algorithm for computing such a triangulation has complexity which is doubly exponential (in k) and produces a doubly exponential sized triangulation, and hence is not suitable for our purpose.

The first main ingredient for our polynomial time algorithm is the double complex associated to the generalized Mayer-Vietoris sequence, which we describe first.

For each S_i, let A_i be the subcomplex corresponding to it, and let $A = A_1 \cap \cdots \cap A_s$. For any $p \geq 0$, let $A^{\alpha_0,\ldots,\alpha_p}$ denote the union, $A_{\alpha_0} \cup \cdots \cup A_{\alpha_p}$. Let $C_i(A)$ denote the $\mathbb{Q}$-vector space of i-chains of A, and $C_\bullet(A)$ the corresponding chain complex.

Proposition 1. The following sequence is exact.

$$0 \longrightarrow C_\bullet(A) \xrightarrow{i} \bigoplus_{\alpha_0} C_\bullet(A^{\alpha_0}) \xrightarrow{\delta} \bigoplus_{\alpha_0<\alpha_1} C_\bullet(A^{\alpha_0,\alpha_1}) \xrightarrow{\delta} \cdots$$

$$\xrightarrow{\delta} \bigoplus_{\alpha_0<\cdots<\alpha_p} C_\bullet(A^{\alpha_0,\ldots,\alpha_p}) \xrightarrow{\delta} \bigoplus_{\alpha_0<\cdots<\alpha_{p+1}} C_\bullet(A^{\alpha_0,\ldots,\alpha_{p+1}}) \xrightarrow{\delta} \cdots,$$

where i is induced by inclusion and the connecting homomorphisms δ are defined as follows:
for $c \in \oplus_{\alpha_0<\cdots<\alpha_p} C_\bullet(A^{\alpha_0,\ldots,\alpha_p})$, $(\delta c)_{\alpha_0,\ldots,\alpha_{p+1}} = \sum_{0\leq i\leq p+1} (-1)^i c_{\alpha_0,\ldots,\hat{\alpha_i},\ldots,\alpha_{p+1}}$.

Proof. See [5]. □

We have a corresponding (fourth quadrant) double complex $\mathcal{N}^{\bullet,\bullet}$, with

$$\mathcal{N}^{p,q} = \oplus_{\alpha_0,\ldots,\alpha_p} C_q(A^{\alpha_0,\ldots,\alpha_p}).$$

$$\begin{array}{ccccccc}
 & & 0 & & 0 & & 0 \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial \\
0 & \longrightarrow & \oplus_{\alpha_0} C_k(A^{\alpha_0}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1} C_k(A^{\alpha_0,\alpha_1}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1<\alpha_2} C_k(A^{\alpha_0,\alpha_1,\alpha_2}) \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial \\
0 & \longrightarrow & \oplus_{\alpha_0} C_{k-1}(A^{\alpha_0}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1} C_{k-1}(A^{\alpha_0,\alpha_1}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1<\alpha_2} C_{k-1}(A^{\alpha_0,\alpha_1,\alpha_2}) \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial \\
0 & \longrightarrow & \oplus_{\alpha_0} C_{k-2}(A^{\alpha_0}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1} C_{k-2}(A^{\alpha_0,\alpha_1}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1<\alpha_2} C_{k-2}(A^{\alpha_0,\alpha_1,\alpha_2}) \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial \\
0 & \longrightarrow & \oplus_{\alpha_0} C_{k-3}(A^{\alpha_0}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1} C_{k-3}(A^{\alpha_0,\alpha_1}) & \xrightarrow{\delta} & \oplus_{\alpha_0<\alpha_1<\alpha_2} C_{k-3}(A^{\alpha_0,\alpha_1,\alpha_2}) \\
 & & \downarrow\partial & & \downarrow\partial & & \downarrow\partial \\
 & & \vdots & & \vdots & & \vdots
\end{array}$$

Standard facts about the Mayer-Vietoris spectral sequence then yields the following well known proposition (see [5]).

Proposition 2. For $0 \leq i \leq k$,

$$H_i(A) \cong H^i(\mathrm{Tot}^\bullet(\mathcal{N}^{\bullet,\bullet})).$$

Note that since $\mathcal{N}^{\bullet,\bullet}$ is a fourth-quadrant double complex, the complex $\mathrm{Tot}^\bullet(\mathcal{N}^{\bullet,\bullet})$ is defined by,

$$\mathrm{Tot}^i(\mathcal{N}^{\bullet,\bullet}) = \bigoplus_{p+k-q=i} \mathcal{N}^{p,q}.$$

Moreover, if we denote by $\mathcal{N}_\ell^{\bullet,\bullet}$ the truncated complex defined by,

$$\begin{array}{rcll}
\mathcal{N}_\ell^{p,q} & = & \mathcal{N}^{p,q}, & 0 \leq p \leq \ell+1,\ \ k-\ell-1 \leq q \leq k, \\
 & = & 0, & \text{otherwise},
\end{array}$$

then it is clear that,

$$H_i(A) \cong H^i(\mathrm{Tot}^\bullet(\mathcal{N}_\ell^{\bullet,\bullet})),\ \ k-\ell \leq i \leq k. \tag{3.1}$$

As noted previously, we cannot hope to compute even the truncated complex $\mathcal{N}_\ell^{\bullet,\bullet}$ since we do not know how to compute triangulations efficiently. We overcome this problem by computing another double complex $\mathcal{D}_\ell^{\bullet,\bullet}$, such that there exists a homomorphism of double complexes,

$$\psi : \mathcal{D}_\ell^{\bullet,\bullet} \to \mathcal{N}_\ell^{\bullet,\bullet},$$

which induces an isomorphism between the E_1 terms of the spectral sequences associated to the double complexes $\mathcal{D}_\ell^{\bullet,\bullet}$ and $\mathcal{N}_\ell^{\bullet,\bullet}$. This implies, by virtue of Theorem 3.1, that the homology groups of the associated total complexes are isomorphic, that is,

$$H^*(\mathrm{Tot}^\bullet(\mathcal{N}_\ell^{\bullet,\bullet})) \cong H^*(\mathrm{Tot}^\bullet(\mathcal{D}_\ell^{\bullet,\bullet})).$$

The construction of the double complex $\mathcal{D}_\ell^{\bullet,\bullet}$ is described in the Section 6.

4. Algorithmic Preliminaries

In this section, we recall some basic algorithms from semi-algebraic geometry that we will need later in the paper. Our main reference is [9]. Here we just describe the inputs, outputs and the complexities of these algorithms.

4.1. **Sign Conditions.** Let R be a real closed field and C its algebraic closure.

A sign condition is an element of $\{0, 1, -1\}$. We define

$$\begin{cases} \mathrm{sign}(x) = 0 & \text{if and only if } x = 0, \\ \mathrm{sign}(x) = 1 & \text{if and only if } x > 0, \\ \mathrm{sign}(x) = -1 & \text{if and only if } x < 0. \end{cases}$$

Let $\mathcal{Q} \subset \mathrm{R}[X_1, \ldots, X_k]$. A sign condition on $\mathcal{Q}$ is an element of $\{0, 1, -1\}^{\mathcal{Q}}$. We say that $\mathcal{Q}$ realizes the sign condition σ at $x \in \mathrm{R}^k$ if

$$\bigwedge_{Q \in \mathcal{Q}} \mathrm{sign}(Q(x)) = \sigma(Q).$$

The realization of the sign condition σ is

$$\mathcal{R}(\sigma) = \{x \in \mathrm{R}^k \mid \bigwedge_{Q \in \mathcal{Q}} \mathrm{sign}(Q(x)) = \sigma(Q)\}.$$

The sign condition σ is realizable if $\mathcal{R}(\sigma)$ is non-empty. We denote by $\mathrm{sign}(\mathcal{Q})$ the set of realizable sign conditions of $\mathcal{Q}$.

4.2. **Representations of points.** Since our algorithms will have to deal with points whose co-ordinates are not rational numbers, we first describe the particular representations of such points that we will use. In the following D denotes an ordered domain contained in the real closed field R.

Let $P \in \mathrm{R}[X]$ and $\sigma \in \{0, 1, -1\}^{\mathrm{Der}(P)}$, a sign condition on the set $\mathrm{Der}(P)$ of derivatives of P. The sign condition σ is a Thom encoding of $x \in \mathrm{R}$ if $\sigma(P) = 0$ and σ is the sign condition taken by the set $\mathrm{Der}(P)$ at x. A real root x of P is completely characterized by the signs of the derivatives of P at x and we say that x is specified by σ. Given a Thom encoding σ, we denote by $x(\sigma)$ the root of P in R specified by σ.

The representations of points are as follows. A k-univariate representation is a $k+2$-tuple of polynomials of $\mathrm{D}[T]$,

$$(f(T), g_0(T), g_1(T), \ldots, g_k(T)),$$

such that f and g_0 are coprime. Note that $g_0(t) \neq 0$ if $t \in \mathrm{C}$ is a root of $f(T)$. The points associated to a univariate representation are the points

$$\left(\frac{g_1(t)}{g_0(t)}, \ldots, \frac{g_k(t)}{g_0(t)}\right) \in \mathrm{C}^k$$

where $t \in \mathrm{C}$ is a root of $f(T)$.

A real k-univariate representation a pair u, σ where u is a k-univariate representation and σ is the Thom encoding of a root of f, $t_\sigma \in \mathrm{R}$. The point associated to the real univariate representation is the point

$$\left(\frac{g_1(t_\sigma)}{g_0(t_\sigma)}, \ldots, \frac{g_k(t_\sigma)}{g_0(t_\sigma)}\right) \in \mathrm{R}^k.$$

4.3. **Complexity of Linear Algebra Subroutines.** In our algorithms, we will also have to compute eigen-vectors of, as well as perform Gram-Schmidt orthogonalizations on, real symmetric matrices whose entries are described by univariate representations. Suppose that M is a matrix of size $k \times k$, given as a real k^2-univariate representation. Recall that this means that M is given as a pair, σ, u, where σ is a Thom encoding and $u = (f(T), g_0(T), g_{1,1}(T), \ldots, g_{k,k}(T))$ is a k^2-univariate representations.

If the degrees of the polynomials $f, g_0, g_{1,1}, \ldots, g_{k,k}$ are bounded by D, then using the most naive algorithms, all the above tasks can be performed with complexity bounded by $(kD)^{O(1)}$. In order to see this notice that the complexity of the algorithms of linear algebra with matrices of size $k \times k$ is bounded by $k^{O(1)}$. However, since the computations takes place in the ring $\mathrm{D}[T]/(f)$, each arithmetic operation in this ring costs $D^{O(1)}$ operations in the ring D.

4.4. **Complexity of Computing Linear Arrangements.** We will also need to compute descriptions of cell complexes, whose cells are the chambers of different dimensions in an arrangement of ℓ hyperplanes in R^k. Notice that the number of cells in such an arrangement can be bounded naively by 3^ℓ. Using standard algorithms, descriptions of the cells in such a complex, and their incidence relationships, can be computed with complexity $k^{O(\ell)}$.

Moreover, in case the arrangement is parametrized, that is the hyperplanes are described by linear equations, whose coeffcients are described by parametrized univariate representations, with at most ℓ parameters, and whose degrees are bounded by D, one can compute a family of polynomials in the parameters, such that for all values of the parameters satisfying a fixed sign condition on this family, the combinatorial structure of the cell complex remain constant. Moreover, the complexity of this procedure, as well as the number and degrees of this new family of polynomials, are all bounded by $(kD)^{O(\ell)}$.

5. Topology of sets defined by quadratic constraints

The results of this section were proved by Agrachev [1] in the context of *non-degenerate* (see [1] for the precise definition of non-degeneracy) quadratic maps. However, for the purposes of this paper the assumption of non-degeneracy is not required as shown below.

Let $P_1, \ldots, P_s$ be homogeneous quadratic polynomials in $\mathrm{R}[X_0, \ldots, X_k]$.

We denote by $P = (P_1, \ldots, P_s) : \mathrm{R}^{k+1} \to \mathrm{R}^s$, the map defined by the polynomials $P_1, \ldots, P_s$.

Let

$$T = \bigcup_{1 \le i \le s} \{x \in \mathbf{S}^k \mid P_i(x) \le 0\}.$$

Let

$$\Omega = \{\omega \in \mathrm{R}^s \mid |\omega| = 1, \omega_i \le 0, 1 \le i \le s\}.$$

For $\omega \in \Omega$ we denote by ωP the quadratic form defined by

$$\omega P = \sum_{i=1}^{s} \omega_i P_i.$$

Let $B \subset \Omega \times \mathbf{S}^k$ be the set defined by,

$$B = \{(\omega, x) \mid \omega \in \Omega, x \in \mathbf{S}^k \text{ and } \omega P(x) \ge 0\}.$$

We denote by $\phi_1 : B \to \Omega$ and $\phi_2 : B \to \mathbf{S}^k$ the two projection maps.

$$\begin{tikzcd} & B \arrow[dl, "\phi_1"'] \arrow[dr, "\phi_2"] & \\ \Omega & & \mathbf{S}^k \end{tikzcd}$$

With the notation developed above,

Proposition 3. The map ϕ_2 gives a homotopy equivalence between B and $\phi_2(B) = T$.

Proof. We first prove that $\phi_2(B) = T$. If $x \in T$, then there exists some $i, 1 \leq i \leq s$, such that $P_i(x) \leq 0$. Then for $\omega = (-\delta_{1,i}, \ldots, -\delta_{s,i})$ (where $\delta_{ij} = 1$ if $i = j$, and 0 otherwise), we see that $(\omega, x) \in B$. Conversely, if $x \in \phi_2(B)$, then there exists $\omega = (\omega_1, \ldots, \omega_s) \in \Omega$ such that, $\sum_{i=1}^{s} \omega_i P_i(x) \geq 0$. Since, $\omega_i \leq 0, 1 \leq i \leq s$, and not all $\omega_i = 0$, this implies that $P_i(x) \leq 0$ for some $i, 1 \leq i \leq s$. This shows that $x \in T$.

For $x \in \phi_2(B)$, the fibre

$$\phi_2^{-1}(x) = \{(\omega, x) \mid \omega \in \Omega \text{ such that } \omega P(x) \geq 0\},$$

is a non-empty subset of Ω defined by a single linear inequality. Thus, each non-empty fiber is an intersection of a convex cone with $\mathbf{S}^s$. Thus, all such fibres can be retracted to their respective centers of mass continuously, proving the first half of the proposition. □

For any quadratic form Q, we will denote by index(Q), the number of negative eigenvalues of the symmetric matrix of the corresponding bilinear form, that is of the matrix M such that, $Q(x) = \langle Mx, x\rangle$ for all $x \in \mathrm{R}^{k+1}$ (here $\langle \cdot, \cdot \rangle$ denotes the usual inner product). We will also denote by $\lambda_i(Q), 0 \leq i \leq k$, the eigenvalues of Q, in non-decreasing order, i.e.

$$\lambda_0(Q) \leq \lambda_1(Q) \leq \cdots \leq \lambda_k(Q).$$

Given a quadratic map $P = (P_1, \ldots, P_s) : \mathrm{R}^{k+1} \to \mathrm{R}^s$, and $0 \leq j \leq k$, we denote by

$$\Omega_j = \{\omega \in \Omega \ \mid \ \lambda_j(\omega P) \geq 0\}.$$

It is clear that the Ω_j's induce a filtration of the space Ω, i.e., $\Omega_0 \subset \Omega_1 \subset \cdots \subset \Omega_k$.

Agrachev [1] showed that the Leray spectral sequence of the map ϕ_1 (converging to the cohomology $H^*(B) \cong H^*(T)$), has as its E_2 terms,

$$E_2^{pq} = H^p(\Omega_{k-q}, \Omega_{k-q-1}). \tag{5.1}$$

(We refer the reader to [12] for the precise definition of the Leray spectral sequence of a map).

Equation 5.1 follows directly from the following lemma.

Lemma 5.1. *The fibre of the map ϕ_1 over a point $\omega \in \Omega_j \setminus \Omega_{j-1}$ has the homotopy type of a sphere of dimension $k - j$.*

Proof. First notice that for $\omega \in \Omega_j \setminus \Omega_{j-1}$, $\lambda_0(\omega P), \ldots, \lambda_{j-1}(\omega P) < 0$. Moreover, letting $Y_0(\omega P), \ldots, Y_k(\omega P)$ be the co-ordinates with respect to an orthonormal basis consisting of the eigenvectors of ωP, we have that $\phi_1^{-1}(\omega)$ is the subset of $\mathbf{S}^k$ defined by,

$$\sum_{i=0}^{k} \lambda_i(\omega P) Y_i(\omega P)^2 \geq 0,$$

$$\sum_{i=0}^{k} Y_i(\omega P)^2 = 1.$$

Since, $\lambda_i(\omega P) < 0, 0 \leq i < j$, it follows that for $\omega \in \Omega_j \setminus \Omega_{j-1}$, the fiber $\phi_1^{-1}(\omega)$ is homotopy equivalent to the $(k-j)$-dimensional sphere defined by setting $Y_0(\omega P) = \cdots = Y_{j-1}(\omega P) = 0$ on the sphere defined by $\sum_{i=0}^{k} Y_i(\omega P)^2 = 1$. □

6. Computing the cohomology groups of a basic semi-algebraic set defined by homogeneous quadratic inequalities

In this section, we will show how to effectively compute the spectral sequence described in the previous section.

Let $\mathcal{P} = (P_1, \ldots, P_s) \subset \mathrm{R}[X_0, \ldots, X_k]$ be a s-tuple of quadratic forms. For any subset $\mathcal{Q} \subset \mathcal{P}$, we denote by $T_{\mathcal{Q}} \subset \mathbf{S}^k$, the semi-algebraic set,

$$T_{\mathcal{Q}} = \bigcup_{P \in \mathcal{Q}} \{x \in \mathbf{S}^k \mid P(x) \leq 0\},$$

and let

$$S = \bigcap_{P \in \mathcal{P}} \{x \in \mathbf{S}^k \mid P(x) \leq 0\}.$$

We denote by $C^\bullet(\mathcal{H}(T_{\mathcal{Q}}))$ the co-chain complex of a cellular subdivision, $\mathcal{H}(T_{\mathcal{Q}})$ of $T_{\mathcal{Q}}$, which is to be chosen sufficiently fine (to be specified later).

We first describe for each subset $\mathcal{Q} \subset \mathcal{P}$ with $\#\mathcal{Q} = \ell < k$, a complex, $\mathcal{M}^\bullet_{\mathcal{Q}}$, and natural homomorphisms,

$$\psi_{\mathcal{Q}} : C^\bullet(\mathcal{H}(T_{\mathcal{Q}})) \to \mathcal{M}^\bullet_{\mathcal{Q}},$$

which induce isomorphisms,

$$\psi^*_{\mathcal{Q}} : H^*(C^\bullet(\mathcal{H}(T_{\mathcal{Q}}))) \to H^*(\mathcal{M}^\bullet_{\mathcal{Q}}).$$

Moreover, for $\mathcal{B} \subset \mathcal{A} \subset \mathcal{P}$ with $\#\mathcal{A} = \#\mathcal{B}+1 < k$, we construct a homomorphism of complexes,

$$\phi_{\mathcal{A},\mathcal{B}} : \mathcal{M}^\bullet_{\mathcal{A}} \to \mathcal{M}^\bullet_{\mathcal{B}},$$

such that the following diagram commutes,

$$\begin{tikzcd} H^*(\mathcal{M}^\bullet_{\mathcal{A}}) \arrow[r, "\phi^*_{\mathcal{A},\mathcal{B}}"] & H^*(\mathcal{M}^\bullet_{\mathcal{B}}) \\ H^*(C^\bullet(\mathcal{H}(T_{\mathcal{A}}))) \arrow[u, "\psi^*_{\mathcal{A}}"] \arrow[r, "r^*"] & H^*(C^\bullet(\mathcal{H}(T_{\mathcal{B}}))) \arrow[u, "\psi^*_{\mathcal{B}}"'] \end{tikzcd} \tag{6.1}$$

where $\phi^*_{\mathcal{A},\mathcal{B}}$ and r^* are the induced homomorphisms of $\phi_{\mathcal{A},\mathcal{B}}$ and the restriction homomorphism r respectively.

Now, consider a fixed subset $\mathcal{Q} \subset \mathcal{P}$, which without loss of generality we take to be $\{P_1, \ldots, P_\ell\}$. Let

$$P_{\mathcal{Q}} = (P_1, \ldots, P_\ell) : \mathrm{R}^{k+1} \to \mathrm{R}^\ell$$

denote the corresponding quadratic map.

As in the previous section, let $\mathrm{R}^{\mathcal{Q}} = \mathrm{R}^\ell$, and

$$\Omega_{\mathcal{Q}} = \{\omega \in \mathrm{R}^\ell \mid |\omega| = 1, \omega_i \leq 0, 1 \leq i \leq \ell\}.$$

Let $B_{\mathcal{Q}} \subset \Omega_{\mathcal{Q}} \times \mathbf{S}^k$ be the set defined by,

$$B_{\mathcal{Q}} = \{(\omega, x) \mid \omega \in \Omega_{\mathcal{Q}}, x \in \mathbf{S}^k \text{ and } \omega P_{\mathcal{Q}}(x) \geq 0\},$$

and we denote by $\phi_{1,\mathcal{Q}} : B_{\mathcal{Q}} \to \Omega_{\mathcal{Q}}$ and $\phi_{2,\mathcal{Q}} : B_{\mathcal{Q}} \to \mathbf{S}^k$ the two projection maps.

For each subset $\mathcal{Q}' \subset \mathcal{Q}$ we have a natural inclusion $\Omega_{\mathcal{Q}'} \hookrightarrow \Omega_{\mathcal{Q}}$.

6.1. **Index Invariant Triangulations.** We now define a certain special kind of semi-algebraic triangulation of $\Omega_{\mathcal{Q}}$ that will play an important role later.

Definition 6.1. (Index Invariant Triangulation) An *index invariant triangulation* of $\Omega_{\mathcal{Q}}$ consists of:

(A) A semi-algebraic triangulation,

$$h_{\mathcal{Q}} : \Delta_{\mathcal{Q}} \to \Omega_{\mathcal{Q}}$$

of $\Omega_{\mathcal{Q}}$, which is compatible with the subsets $\Omega_{\mathcal{Q}'}$ for every $\mathcal{Q}' \subset \mathcal{Q}$, and such that for any simplex σ of $\Delta_{\mathcal{Q}}$, $\mathrm{index}(\omega P_{\mathcal{Q}})$, as well as the multiplicities of the eigenvalues of $\omega P_{\mathcal{Q}}$, stay invariant as ω varies over $h_{\mathcal{Q}}(\sigma)$, and

(B) for every simplex σ of $\Delta_{\mathcal{Q}}$, with $\mathrm{index}(\omega P_{\mathcal{Q}}) = i$, a continuous choice of an orthonormal basis of R^{k+1}, $\{e_0(\sigma, \omega), \ldots, e_k(\sigma, \omega)\}$, for $\omega \in h_{\mathcal{Q}}(\sigma)$, such that $\{e_i(\sigma, \omega), \ldots, e_k(\sigma, \omega)\}$ span the linear subspace of R^{k+1} on which the quadratic form $\omega P_{\mathcal{Q}}$ is positive semi-definite.

We describe later how to compute index invariant triangulations (see Algorithm 1 below). It will follow from the complexity analysis of Algorithm 1 that the size of the complex $\Delta_{\mathcal{Q}}$ as well as the degrees of the polynomials occuring in the parametrized univariate representations defining the various bases, $\{e_0(\sigma, \omega), \ldots, e_k(\sigma, \omega)\}$, are all bounded by $k^{2^{O(\ell)}}$.

For the rest of this section we fix an index invariant triangulation $h_{\mathcal{Q}} : \Delta_{\mathcal{Q}} \to \Omega_{\mathcal{Q}}$, satisfying the complexity estimates stated above.

The following proposition relates, for any simplex $\sigma \in \Delta_{\mathcal{Q}}$, the homotopy type of $\phi^{-1}_{1,\mathcal{Q}}(h_{\mathcal{Q}}(\sigma))$ to that of a single fiber.

Proposition 4. For any simplex $\sigma \in \Delta_{\mathcal{Q}}$ and $\omega \in h_{\mathcal{Q}}(\sigma)$, $\phi^{-1}_{1,\mathcal{Q}}(h_{\mathcal{Q}}(\sigma))$ is homotopy equivalent to $\phi^{-1}_{1,\mathcal{Q}}(\omega)$, and both these spaces have the homotopy type of the sphere $\mathbf{S}^{k-\mathrm{index}(\omega P)}$.

Proof. Let $i = \mathrm{index}(\omega P)$. Since $\mathrm{index}(\omega P)$ is invariant as ω varies over $h_{\mathcal{Q}}(\sigma)$, the quadratic forms ωP has exactly i negative eigen-values for each $\omega \in h_{\mathcal{Q}}(\sigma)$. Let $M(\sigma, \omega) \subset \mathrm{R}^{k+1}$ be the orthogonal complement to the linear span of the corresponding eigen-vectors, and let $B(\sigma, \omega) = M(\sigma, \omega) \cap \mathbf{S}^k$. Clearly, $M(\sigma, \omega)$ and $B(\sigma, \omega)$ vary continuously with ω, and $\phi^{-1}_{1,\mathcal{Q}}(\omega)$ can be retracted to the set $\{\omega\} \times B(\sigma, \omega)$. Finally,

since $h_{\mathcal{Q}}(\sigma)$ is contractible to ω, its clear that $\phi_{1,\mathcal{Q}}^{-1}(h_{\mathcal{Q}}(\sigma))$ retracts to $\{\omega\} \times B(\sigma, \omega)$ and the latter has the homotopy type of $\mathbf{S}^{k-\mathrm{index}(\omega P)}$ by Lemma 5.1. $\square$

6.2. **Definition of the cell complex $\mathcal{K}(B_{\mathcal{Q}})$.** Our next goal is to construct a cell complex homotopy equivalent to $B_{\mathcal{Q}}$ obtained by glueing together certain regular cell complexes, $\mathcal{K}(\sigma)$, where $\sigma \in \Delta_{\mathcal{Q}}$.

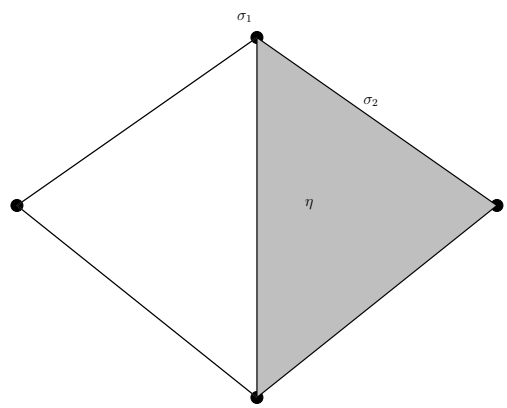

FIGURE 2. The complex $\Delta_{\mathcal{Q}}$.

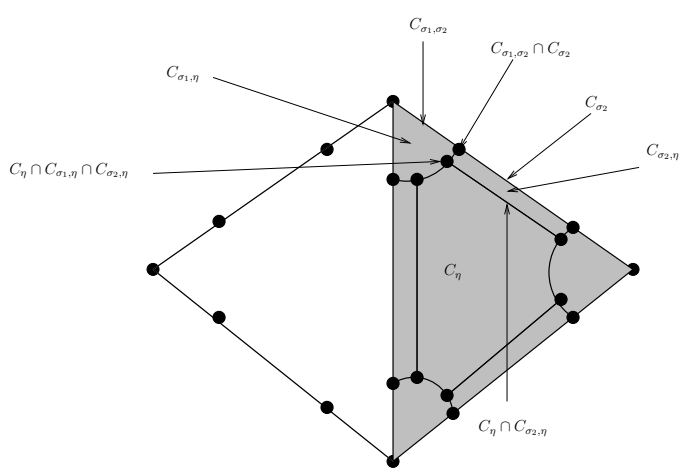

FIGURE 3. The corresponding complex $\mathcal{C}(\Delta_{\mathcal{Q}})$.

Let $1 \geq \varepsilon_0 \gg \varepsilon_1 \gg \cdots \gg \varepsilon_s \gg 0$ be infinitesimals. For $\eta \in \Delta_{\mathcal{Q}}$, we denote by C_η the subset of $\bar{\eta}$ defined by,

$$C_\eta = \{x \in \bar{\eta} \mid \text{ and } \mathrm{dist}(x, \theta) \geq \varepsilon_{\dim(\theta)} \text{ for all } \theta \prec \sigma\}.$$

Now, let $\sigma \prec \eta$ be two simplices of $\Delta_{\mathcal{Q}}$. We denote by $C_{\sigma,\eta}$ the subset of $\bar{\eta}$ defined by,

$$C_{\sigma,\eta} = \{x \in \bar{\eta} \mid \mathrm{dist}(x, \sigma) \leq \varepsilon_{\dim(\sigma)}, \text{ and } \mathrm{dist}(x, \theta) \geq \varepsilon_{\dim(\theta)} \text{ for all } \theta \prec \sigma\}.$$

Note that,

$$|\Delta_{\mathcal{Q}}| = \bigcup_{\sigma \in \Delta_{\mathcal{Q}}} C_\sigma \cup \bigcup_{\sigma,\eta \in \Delta_{\mathcal{Q}}, \sigma \prec \eta} C_{\sigma,\eta}.$$

Also, observe that the various C_η's and $C_{\sigma,\eta}$'s are all homeomorphic to balls, and moreover all non-empty intersections between them also have the same property. Thus, the union of the C_η's and $C_{\sigma,\eta}$'s together with the non-empty intersections between them form a regular cell complex, $\mathcal{C}(\Delta_{\mathcal{Q}})$, whose underlying topological space is $|\Delta_{\mathcal{Q}}|$ (see Figures 2 and 3).

We now associate to each C_σ (respectively, $C_{\sigma,\eta}$) a regular cell complex, $\mathcal{K}(\sigma)$, (respectively, $\mathcal{K}(\sigma,\eta)$) homotopy equivalent to $\phi_{1,\mathcal{Q}}^{-1}(h_{\mathcal{Q}}(C_\sigma))$ (respectively, $\phi_{1,\mathcal{Q}}^{-1}(h_{\mathcal{Q}}(C_{\sigma,\eta}))$.

For each $\sigma \in \Delta_{\mathcal{Q}}$, and $\omega \in h_{\mathcal{Q}}(\sigma)$, let

$$\lambda_0^\sigma(\omega) = \cdots = \lambda_{i_0}^\sigma(\omega) < \lambda_{i_0+1}^\sigma(\omega) = \cdots = \lambda_{i_1}^\sigma(\omega) < \cdots < \lambda_{i_{p-1}+1}^\sigma(\omega) = \cdots = \lambda_{i_p}^\sigma(\omega) < 0 = \cdots = \lambda_{i_{p+1}}^\sigma(\omega) < \cdots = \cdots = \lambda_k^\sigma(\omega)$$

denote the eigenvalues of ωP. Here, $\mathrm{index}(\omega P) = i_p + 1$. Also, since the multiplicities of the eigenvalues do not change as ω varies over $h_{\mathcal{Q}}(\sigma)$, the block structure, $[0,\ldots,i_0],[i_0+1,\ldots,i_1],\ldots,[\cdot,\ldots,k]$ also does not change as ω varies over $h_{\mathcal{Q}}(\sigma)$. For $0 \leq j \leq p$, let $M^j(\sigma,\omega)$ denote the subspace of R^{k+1} orthogonal to the subspace spanned by the eigenvectors corresponding to the eigenvalues $\lambda_0^\sigma(\omega) = \cdots = \lambda_{i_0}^\sigma(\omega) < \lambda_{i_0+1}^\sigma(\omega) = \cdots = \lambda_{i_1}^\sigma(\omega) < \cdots < \lambda_{i_{p-1}+1}^\sigma(\omega) = \cdots = \lambda_{i_j}^\sigma(\omega)$, and let $M(\sigma,\omega) = M^p(\sigma,\omega)$. Since the eigenvalues vary continuously and their multiplicities do not change as ω varies over $h_{\mathcal{Q}}(\sigma)$, the flag of subspaces $M^0(\sigma,\omega) \supset \cdots \supset M^p(\sigma,\omega)$ also varies continuously over $h_{\mathcal{Q}}(\sigma)$.

For each $\sigma \in \Delta_{\mathcal{Q}}$, and $\omega \in h_{\mathcal{Q}}(\sigma)$, let $\{e_0(\sigma,\omega),\ldots,e_k(\sigma,\omega)\}$, be the continuously varying orthonormal basis of R^{k+1} computed previously.

We extend the the orthonormal basis, $\{e_0(\sigma,\omega),\ldots,e_k(\sigma,\omega)\}$, continuously to each $C_{\sigma,\eta}$ for η with $\sigma \prec \eta$, satisfying the condition that

$$M(\sigma,\omega) \subset \mathrm{span}(e_i(\sigma,\omega),\ldots,e_k(\sigma,\omega)).$$

This extension can be done in a consistent manner because the first i eigenvalues, $\lambda_0(\omega P),\ldots,\lambda_{i-1}(\omega P)$ of ωP stay negative, and $\lambda_{i-1}(\omega P) < \lambda_i(\omega P)$ for ω in any infinitesimal neighborhood of $h_{\mathcal{Q}}(C_\sigma)$. Thus, the linear subspace of R^k orthogonal to the eigenspaces corresponding to the eigenvalues, $\lambda_0(\omega P),\ldots,\lambda_{i-1}(\omega P)$ is well defined, and varies continuously with ω in any infinitesimal neighborhood of $h_{\mathcal{Q}}(C_\sigma)$.

More precisely, for any point $z' = h_{\mathcal{Q}}^{-1}(\omega') \in C_{\sigma,\eta}$, let $z = h_{\mathcal{Q}}^{-1}(\omega)$, be the unique point in C_σ closest to z'. For $j = i,\ldots,k$, let $e'_j(\sigma,\omega')$ be the orthogonal projection of $e_j(\sigma,\omega)$ onto the subspace $M(\sigma,\omega')$, and let $e_i(\sigma,\omega'),\ldots,e_k(\sigma,\omega')$ be obtained from $e'_i(\sigma,\omega'),\ldots,e'_k(\sigma,\omega')$ by Gram-Schmidt orthogonalization. Note that each $e_j(\sigma,\omega')$ obtained this way is infinitesimally close to $e_j(\sigma,\omega)$. Using the same procedure starting with the vectors $e_0(\sigma,\omega),\ldots,e_{i-1}(\sigma,\omega)$ and the linear subspace $M(\sigma,\omega')^\perp$, we compute an orthonormal set of vectors $e_0(\sigma,\omega'),\ldots,e_{i-1}(\sigma,\omega')$ spanning $M(\sigma,\omega')^\perp$. It is clear from construction, that the vectors $e_0(\sigma,\omega'),\ldots,e_k(\sigma,\omega')$ form an orthonormal basis for R^{k+1}. Moreover, they vary continuously with ω', extending the basis $e_0(\sigma,\omega),\ldots,e_k(\sigma,\omega)$, and finally $e_i(\sigma,\omega'),\ldots,e_k(\sigma,\omega')$ span $M(\sigma,\omega')$.

The orthonormal basis

$$\{e_0(\sigma,\omega),\ldots,e_k(\sigma,\omega)\},$$

determines a complete flag of subspaces, $\mathcal{F}(\sigma,\omega)$, consisting of

$$\begin{aligned} L^0(\sigma,\omega) &= 0,\\ L^1(\sigma,\omega) &= \mathrm{span}(e_k(\sigma,\omega)),\\ L^2(\sigma,\omega) &= \mathrm{span}(e_k(\sigma,\omega),e_{k-1}(\sigma,\omega)),\\ &\vdots\\ L^{k+1}(\sigma,\omega) &= \mathrm{R}^{k+1}. \end{aligned}$$

For $0 \le j \le k$, let $c_j^+(\sigma,\omega)$ (respectively, $c_j^-(\sigma,\omega)$) denote the $(k-j)$-dimensional cell consisting of the intersection of the $L^{k-j+1}(\sigma,\omega)$ with the unit hemisphere in R^{k+1} defined by $\{x \in \mathbf{S}^k \mid \langle x, e_j(\sigma,\omega)\rangle \ge 0\}$ (respectively, $\{x \in \mathbf{S}^k \mid \langle x, e_j(\sigma,\omega)\rangle \le 0\}$).

The regular cell complex $\mathcal{K}(\sigma)$ (as well as $\mathcal{K}(\sigma,\eta)$) is defined as follows.

The cells of $\mathcal{K}(\sigma)$ are $\{(x,\omega) \mid x \in c_j^{\pm}(\sigma,\omega), \omega \in h_{\mathcal{Q}}(c)\}$, where $\mathrm{index}(\omega P) \le j \le k$, and $c \in \mathcal{C}(\Delta_{\mathcal{Q}})$ is either C_σ itself, or a cell contained in the boundary of C_σ.

Similarly, the cells of $\mathcal{K}(\sigma,\eta)$ are $\{(x,\omega) \mid x \in c_j^{\pm}(\sigma,\omega), \omega \in h_{\mathcal{Q}}(c)\}$, where $\mathrm{index}(\omega P) \le j \le k$, $c \in \mathcal{C}(\Delta_{\mathcal{Q}})$ is either $C_{\sigma,\eta}$ itself, or a cell contained in the boundary of $C_{\sigma,\eta}$.

Our next step is to obtain cellular subdivisions of each non-empty intersection amongst the spaces associated to the complexes constructed above, and thus obtain a regular cell complex, $\mathcal{K}(B_{\mathcal{Q}})$, whose associated space, $|\mathcal{K}(B_{\mathcal{Q}})|$, will be shown to be homotopy equivalent to $B_{\mathcal{Q}}$ (Proposition 6 below).

First notice that $|\mathcal{K}(\sigma',\eta')|$ (respectively, $|\mathcal{K}(\sigma)|$) has a non-empty intersection with $|\mathcal{K}(\sigma,\eta)|$ only if $C_{\sigma',\eta'}$ (respectively, $C_{\sigma'}$) intersects $C_{\sigma,\eta}$.

Let C be some non-empty intersection amongst the C_σ's and $C_{\sigma,\eta}$'s, that is C is a cell of $\mathcal{C}(\Delta_{\mathcal{Q}})$. Then, $C \subset \eta$ for a unique simplex $\eta \in \Delta_{\mathcal{Q}}$, and

$$C = C_{\sigma_1,\eta} \cap \cdots \cap C_{\sigma_p,\eta},$$

with $\sigma_1 \prec \sigma_2 \prec \cdots \prec \sigma_p \prec \eta$ and $p \le \#\mathcal{Q}+1$.

Consider $\omega \in h_{\mathcal{Q}}(C)$. We have p different flags,

$$\mathcal{F}(\sigma_1,\omega),\ldots,\mathcal{F}(\sigma_p,\omega),$$

giving rise to p independent regular cell decompositions of $B(\omega,\eta) = M(\eta,\omega) \cap \mathbf{S}^k$.

There is a unique smallest regular cell complex, $\mathcal{K}'(C,\omega)$, that refines all these cell decompositions. The cells of this cell decomposition consists of the following. Let $L \subset M(\eta,\omega)$ be any linear subspace of dimension $m, 0 \le m \le k+1$, which is an intersection, of linear subspaces $L_1,\ldots,L_p$, where $L_i \in \mathcal{F}(\sigma_i,\omega), 1 \le i \le p$. The elements of the flags, $\mathcal{F}(\sigma_1,\omega),\ldots,\mathcal{F}(\sigma_p,\omega)$ of dimensions $m+1$, partition L into polyhedral cones of various dimensions. The union of the sets of intersections of these cones with $\mathbf{S}^k$, over all such subspaces $L \subset M(\eta,\omega)$, are the cells of $\mathcal{K}'(C,\omega)$. Figure 4 illustrates the refinement described above in case of two flags in R^3.

We now triangulate $h_{\mathcal{Q}}(C)$, using the algorithm implicit in Theorem 3.2 (Triangulation), such that the combinatorial type of the arrangement of flags,

$$\mathcal{F}(\sigma_1,\omega),\ldots,\mathcal{F}(\sigma_p,\omega)$$

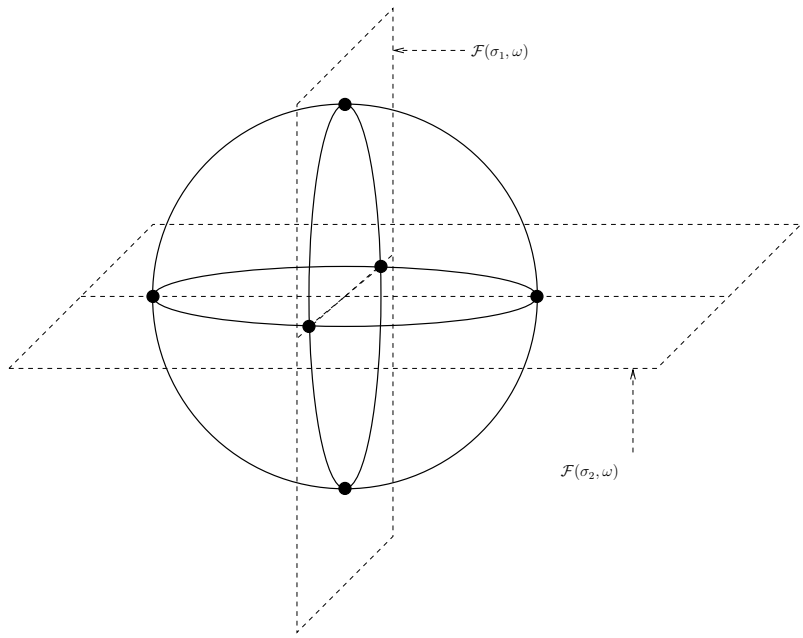

FIGURE 4. The cell complex $\mathcal{K}'(C, \omega)$.

and hence the cell decomposition $\mathcal{K}'(C, \omega)$, stays invariant over the image, $h_C(\theta)$, of each simplex, θ, of this triangulation. More precisely, we first compute a family of polynomials, $\mathcal{A}_C \subset \mathrm{R}[Z_1, \ldots, Z_\ell]$ whose signs at ω determine the combinatorial type of the corresponding arrangement of flags. It is easy to verify (see Section 4.4), given the complexity bounds on the parametrized univariate representations defining the orthonormal bases, $\{e_0(\sigma, \omega), \ldots, e_k(\sigma, \omega)\}$, $\omega \in h_{\mathcal{Q}}(\sigma)$, stated above, that the number and degrees of the polynomials in the family $\mathcal{A}_C$ is bounded by $k^{2^{O(\ell)}}$. We then use the algorithm implicit in Theorem 3.2 (Triangulation), with $\mathcal{A}_C$ as input, to obtain the required triangulation.

The closures of the sets

$$\{(\omega, x) \mid x \in c \in \mathcal{K}'(C, \omega), \omega \in h_{\mathcal{Q}}(h_C(\theta))\}$$

constitute a regular cell complex, $\mathcal{K}(C)$, which is compatible with the regular cell complexes $\mathcal{K}(\sigma_1), \ldots, \mathcal{K}(\sigma_p)$.

The following proposition gives an upper bound on the size of the complex $\mathcal{K}(C)$. We use the notation introduced in the previous paragraph.

Proposition 5. For each $\omega \in h_{\mathcal{Q}}(C)$, the number of cells in $\mathcal{K}'(C, \omega)$ is bounded by $k^{O(\ell)}$. Moreover, the number of cells in the complex $\mathcal{K}(C)$ is bounded by $k^{2^{O(\ell)}}$.

Proof. The first part of the proposition follows from the fact that there are at most $k^{\#\mathcal{Q}+1} = k^{\ell+1}$ choices for the linear space L and the number of $(m-1)$ dimensional cells contained in L is bounded by 2^ℓ (which is an upper bound on the number of full dimensional cells in an arrangement of at most ℓ hyperplanes). The second part is a consequence of the complexity estimate in Theorem 3.2 (Triangulation) and the bounds on number and degrees of polynomials in the family $\mathcal{A}_C$ stated above. $\square$

We denote by $\mathcal{K}(B_{\mathcal{Q}})$, the union of all the complexes $\mathcal{K}(C)$ constructed above, noting that by construction, $\mathcal{K}(B_{\mathcal{Q}})$ is a regular cell complex.

Proposition 6. $|\mathcal{K}(B_{\mathcal{Q}})|$ is homotopy equivalent to $B_{\mathcal{Q}}$.

Proof. We first show that $B_{\mathcal{Q}}$ is homotopy equivalent to a subset $B'_{\mathcal{Q}} \subset B_{\mathcal{Q}}$ as follows. For each simplex $\sigma \in \Delta_{\mathcal{Q}}$ of the largest dimension ℓ, we use the retraction used in the proof of Proposition 4, to retract $\phi_1^{-1}(h_{\mathcal{Q}}(\mathrm{relint}(C_\sigma)))$ to the set $\{(\omega, x) \mid \omega \in \mathrm{relint}(C_\sigma), x \in B(\omega, \sigma)\}$. In this way we obtain a semi-algebraic set, X'_ℓ, which is a deformation retract of $\mathrm{Ext}(B_{\mathcal{Q}}, \mathrm{R}\langle \varepsilon_0, \ldots, \varepsilon_{\ell-1}\rangle)$. Let $X_\ell = \lim_{\varepsilon_{\ell-1}} X'_\ell$. Notice that in the definition of X'_ℓ, if we replace $\varepsilon_{\ell-1}$ by a variable t and denote the corresponding set by $X'_{\ell,t}$, then for all $0 < t < t'$, $X'_{\ell,t} \subset X'_{\ell,t'}$ and each $X_{\ell,t}$ is closed and bounded. It then follows (see Lemma 16.17 in [9]) that $\mathrm{Ext}(X_\ell, \mathrm{R}\langle \varepsilon_0, \ldots, \varepsilon_{\ell-1}\rangle)$ has the same homotopy type as X'_ℓ, and hence X_ℓ has the same homotopy type as $\mathrm{Ext}(B_{\mathcal{Q}}, \mathrm{R}\langle \varepsilon_0, \ldots, \varepsilon_{\ell-2}\rangle)$.

Now repeat the process using the $(\ell-1)$-dimensional simplices and so on, to finally obtain $X_0 = B'_{\mathcal{Q}}$, which by construction has the same homotopy type as $B_{\mathcal{Q}}$. Finally, (again using Lemma 16.17 in [9]) we also have that $X_0 = \lim_{\varepsilon_0} |\mathcal{K}(B_{\mathcal{Q}})|$ and $\mathrm{Ext}(X_0, \mathrm{R}\langle \varepsilon_0, \ldots, \varepsilon_{\ell-1}\rangle)$ has the same homotopy type as $|\mathcal{K}(B_{\mathcal{Q}})|$. □

We also have,

Proposition 7. The number of cells in the cell complex $\mathcal{K}(B_{\mathcal{Q}})$ is bounded by $k^{2^{O(\ell)}}$.

Proof. The proposition is a consequence of Proposition 5 and the fact that the number of cells in the complex $\mathcal{C}(\Delta_{\mathcal{Q}})$ is bounded by $k^{2^{O(\ell)}}$. □

We now define,

$$\mathcal{M}^\bullet_{\mathcal{Q}} = C^\bullet(\mathcal{K}(B_{\mathcal{Q}}),$$

where $C^\bullet(\mathcal{K}(B_{\mathcal{Q}})$ is the cellular co-chain complex of the regular cell complex $\mathcal{K}(B_{\mathcal{Q}})$.

Let $\mathcal{H}(T_{\mathcal{Q}})$ (resp. $\mathcal{H}(B_{\mathcal{Q}})$) be a suitably fine cellular subdivision of $T_{\mathcal{Q}}$ (resp. $B_{\mathcal{Q}}$) and let

$$\phi'_{2,\mathcal{Q}} : C_\bullet(\mathcal{H}(B_{\mathcal{Q}})) \to C_\bullet(\mathcal{H}(T_{\mathcal{Q}})),$$

be the homomorphism induced by a cellular map, which is a cellular approximation of $\phi_{2,\mathcal{Q}}$.

Let $\phi_{\mathcal{Q}} : |\mathcal{K}(B_{\mathcal{Q}})| \to B_{\mathcal{Q}}$ denote the homotopy equivalence shown to exist by Proposition 6 above and let

$$\phi'_{\mathcal{Q}} : C_\bullet(\mathcal{K}'(B_{\mathcal{Q}})) \to C_\bullet(\mathcal{H}(B_{\mathcal{Q}})),$$

be the homomorphism induced by a cellular approximation to $\phi_{\mathcal{Q}}$, where $\mathcal{K}'(B_{\mathcal{Q}})$ is a cellular subdivision of the complex $\mathcal{K}(B_{\mathcal{Q}})$.

Since, each cell of $\mathcal{K}(B_{\mathcal{Q}})$ is a union of cells of $\mathcal{K}'(B_{\mathcal{Q}})$, there is a natural homomorphism

$$\theta_{\mathcal{Q}} : C_\bullet(\mathcal{K}(B_{\mathcal{Q}})) \to C_\bullet(\mathcal{K}'(B_{\mathcal{Q}}))$$

obtained by sending each p-dimensional cell of $\mathcal{K}(B_{\mathcal{Q}})$ to the sum of p-dimensional cells of $\mathcal{K}'(B_{\mathcal{Q}})$ contained in it, for every $p \geq 0$. It is a standard fact that $\theta_{\mathcal{Q}}$ and its dual, $\check{\theta}_{\mathcal{Q}}$, are quasi-isomorphisms.

Let

$$\psi_{\mathcal{Q}} = \check{\theta}_{\mathcal{Q}} \circ \check{\phi}'_{\mathcal{Q}} \circ \check{\phi}'_{2,\mathcal{Q}} : C^\bullet(\mathcal{H}(T_{\mathcal{Q}})) \to C^\bullet(\mathcal{K}(B_{\mathcal{Q}})),$$

where $\check{\phi}'_{\mathcal{Q}}$ (resp. $\check{\phi}'_{2,\mathcal{Q}}$) is the dual homomorphism of $\phi'_{\mathcal{Q}}$ (resp. $\phi'_{2,\mathcal{Q}}$).

Proposition 8. For $0 \leq i \leq k-1$, the induced homomorphisms,

$$\psi^*_{\mathcal{Q}} : H^i(C^\bullet(\mathcal{H}(T_{\mathcal{Q}}))) \to H^i(\mathcal{M}^\bullet_{\mathcal{Q}})$$

are isomorphisms.

Proof. The proof is clear since $\psi_{\mathcal{Q}}$ is a composition of quasi-isomorphisms. □

Now let, $\mathcal{B} \subset \mathcal{A} \subset \mathcal{P}$ with $\#\mathcal{A} = \#\mathcal{B} + 1 < k$.

The simplicial complex $\Delta_{\mathcal{B}}$ is a subcomplex of $\Delta_{\mathcal{A}}$ and hence, $\mathcal{K}(B_{\mathcal{B}})$ is a subcomplex of $\mathcal{K}(B_{\mathcal{A}})$ and thus there exists a natural homomorphism (induced by restriction),

$$\phi_{\mathcal{A},\mathcal{B}} : \mathcal{M}^\bullet_{\mathcal{A}} \to \mathcal{M}^\bullet_{\mathcal{B}}.$$

The complexes $\mathcal{M}^\bullet_{\mathcal{A}}, \mathcal{M}^\bullet_{\mathcal{B}}$, and the homomorphisms, $\phi_{\mathcal{A},\mathcal{B}}, \psi_{\mathcal{A}}, \psi_{\mathcal{B}}$ satisfy

Proposition 9. The diagram

$$\begin{array}{ccc} \mathcal{M}^\bullet_{\mathcal{A}} & \xrightarrow{\phi_{\mathcal{A},\mathcal{B}}} & \mathcal{M}^\bullet_{\mathcal{B}} \\ \uparrow \psi_{\mathcal{A}} & & \uparrow \psi_{\mathcal{B}} \\ C^\bullet(\mathcal{H}(T_{\mathcal{A}})) & \xrightarrow{r} & C^\bullet(\mathcal{H}(T_{\mathcal{B}})) \end{array}$$

is commutative, where r is the restriction homomorphism.

Proof. Clear from the construction. □

It follows from Proposition 9 that the diagram (6.1) is also commutative.

We denote by

$$\check{\phi}_{\mathcal{B},\mathcal{A}} : \check{\mathcal{M}}^\bullet_{\mathcal{B}} \to \check{\mathcal{M}}^\bullet_{\mathcal{A}}$$

the homomorphism dual to $\phi_{\mathcal{A},\mathcal{B}}$. We denote by $\mathcal{D}^{\bullet,\bullet}_{\mathcal{P}}$ the double complex defined by:

$$\mathcal{D}^{p,q}_{\mathcal{P}} = \bigoplus_{\mathcal{Q} \subset \mathcal{P}, \#\mathcal{Q} = p+1} \check{\mathcal{M}}^q_{\mathcal{Q}}.$$

The vertical differentials,

$$d : \mathcal{D}^{p,q}_{\mathcal{P}} \to \mathcal{D}^{p,q-1}_{\mathcal{P}},$$

are induced componentwise from the differentials of the individual complexes $\check{\mathcal{M}}^\bullet_{\mathcal{Q}}$. The horizontal differentials,

$$\delta : \mathcal{D}^{p,q}_{\mathcal{P}} \to \mathcal{D}^{p+1,q}_{\mathcal{P}},$$

are defined as follows: for $a \in \mathcal{D}^{p,q}_{\mathcal{P}} = \oplus_{\#\mathcal{Q}=p+1} \check{\mathcal{M}}^q_{\mathcal{Q}}$, and for each subset

$$\mathcal{Q} = \{P_{i_0}, \ldots, P_{i_{p+1}}\} \subset \mathcal{P}$$

with $i_0 < \cdots < i_{p+1}$, the $\mathcal{Q}$-th component of $\delta a \in \mathcal{D}^{p+1,q}_{\mathcal{P}}$ is given by,

$$(\delta a)_{\mathcal{Q}} = \sum_{0 \leq j \leq p+1} \check{\phi}_{\mathcal{Q}_j,\mathcal{Q}}(a_{\mathcal{Q}_j}),$$

where $\mathcal{Q}_j = \mathcal{Q} \setminus \{P_{i_j}\}$.

$$\begin{array}{ccccccccc}
 & & \vdots & & \vdots & & \vdots & & \\
 & & \downarrow d & & \downarrow d & & \downarrow d & & \\
0 & \longrightarrow & \oplus_{\#\mathcal{Q}=1}\check{\mathcal{M}}^3_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=2}\check{\mathcal{M}}^3_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=3}\check{\mathcal{M}}^3_{\mathcal{Q}} & \longrightarrow & \cdots \\
 & & \downarrow d & & \downarrow d & & \downarrow d & & \\
0 & \longrightarrow & \oplus_{\#\mathcal{Q}=1}\check{\mathcal{M}}^2_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=2}\check{\mathcal{M}}^2_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=3}\check{\mathcal{M}}^2_{\mathcal{Q}} & \longrightarrow & \cdots \\
 & & \downarrow d & & \downarrow d & & \downarrow d & & \\
0 & \longrightarrow & \oplus_{\#\mathcal{Q}=1}\check{\mathcal{M}}^1_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=2}\check{\mathcal{M}}^1_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=3}\check{\mathcal{M}}^1_{\mathcal{Q}} & \longrightarrow & \cdots \\
 & & \downarrow d & & \downarrow d & & \downarrow d & & \\
0 & \longrightarrow & \oplus_{\#\mathcal{Q}=1}\check{\mathcal{M}}^0_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=2}\check{\mathcal{M}}^0_{\mathcal{Q}} & \xrightarrow{\delta} & \oplus_{\#\mathcal{Q}=3}\check{\mathcal{M}}^0_{\mathcal{Q}} & \longrightarrow & \cdots \\
 & & \downarrow d & & \downarrow d & & \downarrow d & & \\
 & & 0 & & 0 & & 0 & &
\end{array}$$

We have the following theorem.

Theorem 6.2. *For* $0 \leq i \leq k$,

$$H^i(S) \cong H^i(\mathrm{Tot}^\bullet(\mathcal{D}^{\bullet,\bullet}_{\mathcal{P}})).$$

Proof. For each $i, 1 \leq i \leq s$, let $S_i \subset \mathbf{S}^k$ denote, the set defined on $\mathbf{S}^k$ by $P_i \leq 0$. Then, $S = \cap_{i=1}^s S_i$. Choosing a suitably fine triangulation of $\cup_{i=1}^s S_i$ consider the Mayer-Vietoris double complex, $\mathcal{N}^{\bullet,\bullet}$, as described in Section 3.5. The homomorphisms,

$$\bigoplus_{\mathcal{Q}\subset\mathcal{P},\#\mathcal{Q}=p+1} \check{\psi}_{\mathcal{Q}} : \bigoplus_{\mathcal{Q}\subset\mathcal{P},\#\mathcal{Q}=p+1} \check{\mathcal{M}}^q_{\mathcal{Q}} \longrightarrow \bigoplus_{\mathcal{Q}\subset\mathcal{P},\#\mathcal{Q}=p+1} C_q(\mathcal{H}(T_{\mathcal{Q}}))$$

give a homomorphism of the double complexes,

$$\psi : \mathcal{D}^{\bullet,\bullet}_{\mathcal{P}} \longrightarrow \mathcal{N}^{\bullet,\bullet}. \tag{6.2}$$

By Proposition 8, ψ induces isomporphisms between the E_1 terms of the two spectral sequences, obtained by taking homology with respect to the vertical differentials. Theorem 3.1 then implies that ψ induces isomorphisms in the associated spectral sequences. But, the second spectral sequence converges to the homology of S by (3.1). The theorem is an immediate consequence. □

7. Algorithms for quadratic forms

In this section we describe the algorithm for computing the top Betti numbers of a basic semi-algebraic set defined by quadratic forms. We first describe an algorithm for computing index invariant triangulations (see Definition 6.1).

Algorithm 1 (Index Invariant Triangulation).

Input A set $\mathcal{Q} = \{P_1, \ldots, P_\ell\} \subset \mathrm{R}[X_0, \ldots, X_k]$ where each P_i is a quadratic form. We denote by $P_{\mathcal{Q}} = (P_1, \ldots, P_\ell) : \mathrm{R}^{k+1} \to \mathrm{R}^\ell$ the corresponding quadratic map.

Output

(A) A semi-algebraic triangulation,

$$h_{\mathcal{Q}} : \Delta_{\mathcal{Q}} \to \Omega_{\mathcal{Q}}$$

of $\Omega_{\mathcal{Q}}$, which is compatible with the subsets $\Omega_{\mathcal{Q}'}$ for every $\mathcal{Q}' \subset \mathcal{Q}$, and such that for any simplex σ of $\Delta_{\mathcal{Q}}$, index(ωP), as well as the multiplicities of the eigenvalues of ωP, stay invariant as ω varies over $h_{\mathcal{Q}}(\sigma)$.

(B) For each simplex σ of $\Delta_{\mathcal{Q}}$, a set of parametrized univariate representations, $\{u_0(\sigma,\omega),\ldots,u_k(\sigma,\omega)\}$, parametrized by $\omega \in h_{\mathcal{Q}}(\sigma)$, such that the associated points in R^{k+1},

$$\{e_0(\sigma,\omega),\ldots,e_k(\sigma,\omega)\},$$

form an orthornormal basis of R^{k+1}, and

$$\{e_i(\sigma,\omega),\ldots,e_k(\sigma,\omega)\}$$

span the linear subspace of R^{k+1} on which the quadratic form $\omega P_{\mathcal{Q}}$ is positive semi-definite (here $i = \mathrm{index}(\omega P_{\mathcal{Q}})$).

Procedure

Step 1 Let $\varepsilon > 0$ be an infinitesimal and let $Z = (Z_1,\ldots,Z_\ell)$. Also, let $M_{\mathcal{Q}}$ be the symmetric matrix corresponding to the quadratic form (in $X_0,\ldots,X_k$) defined by

$$\overline{P}_{\mathcal{Q}}(Z) = (1-\varepsilon)(Z_1P_1 + \cdots + Z_\ell P_\ell) + \varepsilon Q,$$

where $Q = \sum_{i=0}^{k} iX_i^2$. The entries of $M_{\mathcal{Q}}$ depend linearly on $Z_1,\ldots,Z_\ell$, and ε. Compute the polynomial,

$$F(Z,T) = \det(M_{\mathcal{Q}} + T\cdot I_{k+1}) = T^{k+1} + C_kT^k + \cdots + C_0,$$

where each $C_i \in \mathrm{R}[\varepsilon][Z_1,\ldots,Z_\ell]$ is a polynomial of degree at most $k+1$.

Step 2 Using Algorithm 11.1 in [9] (Elimination), compute a family of polynomials $\mathcal{A}_{\mathcal{Q}} \subset \mathrm{R}[Z_1,\ldots,Z_\ell]$ such that such that for each $\rho \in \mathrm{Sign}(\mathcal{A}_{\mathcal{Q}})$, and $z \in \mathcal{R}(\rho)$, the Thom encodings of the roots of $F(z,T)$ in $\mathrm{R}\langle\varepsilon\rangle$, as well as the the number of non-negative roots of $F(z,T)$ stay constant.

Step 3 Using the algorithm implicit in Theorem 3.2 (Triangulation), compute a semi-algebraic triangulation,

$$h_{\mathcal{Q}} : \Delta_{\mathcal{Q}} \to \Omega_{\mathcal{Q}},$$

respecting the family $\mathcal{A}_{\mathcal{Q}} \cup \cup_{i=1}^{\ell}\{Z_i\}$.

Step 4 For each simplex σ of $\Delta_{\mathcal{Q}}$, let $\tau_0(\sigma),\ldots,\tau_k(\sigma)$ be the Thom encodings of the real roots, $\lambda_0(\sigma,z) < \cdots < \lambda_k(\sigma,z)$ of $F(z,T)$, for $z \in h_{\mathcal{Q}}(\sigma)$. For $0 \le i \le k$, compute using linear algebra a parametrized univariate representation

$$\bar{u}_i(Z,T) = (F, g_{i,0},\ldots,g_{i,k+1}),$$

such that for each $z \in h_{\mathcal{Q}}(\sigma)$ the point associated to $\bar{u}_i(z,T)$ is the eigenvector of $M_{\mathcal{Q}}(z)$ corresponding to the eigenvalue $\lambda_i(z)$. Let $u_i = \lim_\varepsilon \bar{u}_i$, and let $e_i(z) \in \mathrm{R}^{k+1}$ denote the corresponding unit vector.

Complexity Analysis: The complexity of the algorithm is dominated by the complexity of Step 3, which is $k^{2^{O(\ell)}}$. $\square$

Proof of Correctness: The quadratic form $\overline{P}_{\mathcal{Q}}(z)$, and hence the matrix $M_{\mathcal{Q}}$, has $k+1$ distinct eigen values for each $z \in \mathrm{R}^\ell$ and has the same index as $zP_{\mathcal{Q}}$. To prove the first part, replace ε in the definition of $\overline{P}_{\mathcal{Q}}(Z)$ and observe that the statement is true $t = 1$, since the quadratic form Q has distinct eigenvalues. Thus, the set of t's for which $\overline{P}_{\mathcal{Q}}(z)$ has $k+1$ distinct eigenvalues is non-empty,

constructible and contains a open subset, since the condition of having distinct eigenvalues is a stable condition. Thus, there exists $\varepsilon_0 > 0$, such that for all $t \in (0, \varepsilon_0)$, $\overline{P}_{\mathcal{Q}}(z)$ has $k+1$ distinct eigen-values, and hence it is also the case for any infinitesimal t. The fact that $\overline{P}_{\mathcal{Q}}(z)$ has the same index as $zP_{\mathcal{Q}}$ follows from the fact that the quadratic form Q is positive semi-definite. Thus, the eigenspaces corresponding to the eigenvalues of $M_{\mathcal{Q}}(z)$ are all one-dimensional and thus the vectors $e'_i(\sigma, \omega)$ all well defined (upto multiplication by -1). Finally, note that since $e'_0(\sigma, \omega), \ldots, e'_k(\sigma, \omega)$ are orthonormal, so are $e_0(\sigma, \omega), \ldots, e_k(\sigma, \omega)$ for every $\omega \in h_{\mathcal{Q}}(\sigma)$, and since $e'_i(\sigma, \omega), \ldots, e'_k(\sigma, \omega)$ span the non-negative eigenspace of $\omega\overline{P}_\omega$, their images under the $\lim_\varepsilon$ map will span the non-negative eigenspace of $\omega P_{\mathcal{Q}}$. $\square$

We now describe an algorithm for computing the complexes $\check{\mathcal{M}}^\bullet_{\mathcal{Q}}$ described in the previous section.

Algorithm 2 (Build Complex for Unions).

Input (A) An integer $\ell, 0 \leq \ell \leq k$.
(B) A quadratic map $P = (P_1, \ldots, P_s) : \mathrm{R}^{k+1} \to \mathrm{R}^s$ given by s homogeneous quadratic polynomials, $P_1, \ldots, P_s \in \mathrm{R}[X_0, \ldots, X_k]$.

Output (A) For each subset $\mathcal{Q} \subset \mathcal{P} = \{P_1, \ldots, P_s\}, \#\mathcal{Q} \leq \ell + 2$ a description of the complex $\check{\mathcal{M}}_{\mathcal{Q}}$, consisting of a basis for each term of the complex and matrices (in this basis) for the differentials.
(B) For each $\mathcal{Q}' \subset \mathcal{Q}$, with $\#\mathcal{Q} = \#\mathcal{Q}' + 1$, matrices for the homomorphisms,

$$\check{\phi}_{\mathcal{Q}',\mathcal{Q}} : \check{\mathcal{M}}^\bullet_{\mathcal{Q}'} \to \check{\mathcal{M}}^\bullet_{\mathcal{Q}}.$$

Procedure

Step 1 For each subset $\mathcal{Q} = \{P_{i_1}, \ldots, P_{i_{\ell+2}}\} \subset \mathcal{P}$, with $\#\mathcal{Q} = \ell + 2$, let $P_{\mathcal{Q}}$ be the quadratic map corresponding to the subset $\mathcal{Q}$. Call Algorithm 1(Index Invariant Triangulation) with input $\mathcal{Q}$.

Step 2 Construct the cell complex $\mathcal{C}(\Delta_{\mathcal{Q}})$ (following its definition given in Section 6.2).

Step 3 For each cell $C \in \mathcal{C}(\Delta_{\mathcal{Q}})$, compute using the algorithm implicit in Theorem 3.2 (Triangulation), the cell complex $\mathcal{K}(C)$ and thus obtain a description of $\mathcal{K}(B_{\mathcal{Q}})$.

Step 4 Compute the matrices corresponding to the differentials in the complex $\mathcal{M}^\bullet_{\mathcal{Q}} = C^\bullet(\mathcal{K}(B_{\mathcal{Q}}))$.

Step 5 For $\mathcal{Q}' \subset \mathcal{Q} \subset \mathcal{P}$ with $\#\mathcal{Q} = \#\mathcal{Q}' + 1 < k$, compute the matrices for the homomorphisms of complexes,

$$\check{\phi}_{\mathcal{Q}',\mathcal{Q}} : \check{\mathcal{M}}^\bullet_{\mathcal{Q}'} \to \check{\mathcal{M}}^\bullet_{\mathcal{Q}}.$$

in the following way.

The simplicial complex $\mathcal{K}(B_{\mathcal{Q}'})$ is a subcomplex of $\mathcal{K}(B_{\mathcal{Q}})$ by construction. Compute the matrix for the restriction homomorphism,

$$\phi_{\mathcal{Q},\mathcal{Q}'} : C^\bullet(\mathcal{K}(B_{\mathcal{Q}})) \to C^\bullet(\mathcal{K}(B_{\mathcal{Q}'})).$$

and output the matrix for the dual homomorphism.

Complexity Analysis: The complexity of Step 1 $\sum_{i=0}^{\ell+2} \binom{s}{i} k^{O(i)}$, using the complexity of Algorithm 1. The complexity of Step 2 is $\sum_{i=0}^{\ell+2} \binom{s}{i} k^{2^{O(\min(\ell,s))}}$, using the complexity of the algorithm for triangulating semi-algebraic sets. It follows from

Proposition 7 that the complexities of all the remaining steps are also bounded by $\sum_{i=0}^{\ell+2}\binom{s}{i}k^{2^{O(\min(\ell,s))}}$. □

Proof of Correctness: The correctness of the algorithm is a consequence of the correctness of Algorithm 1 and Propositions 6 and 9. □

Let $P_1,\ldots,P_s \in \mathrm{R}[X_0,\ldots,X_k]$ be homogeneous quadratic polynomials, and consider the set $S \subset \mathbf{S}^k$ defined by, $S = \{x \in \mathbf{S}^k \mid P_1(x) \leq 0,\ldots,P_s \leq 0\}$.

We will also denote for $1 \leq i \leq s$, by S_i the set defined by $\{x \in \mathbf{S}^k \mid P_i(x) \leq 0.\}$

Clearly, $S = \cap_{1\leq i\leq s} S_i$.

Algorithm 3 (Computing the highest ℓ Betti Numbers: the homogeneous case).

Input A quadratic map $P = (P_1,\ldots,P_s) : \mathrm{R}^{k+1} \to \mathrm{R}^s$ given by a set, $\mathcal{P} = \{P_1,\ldots,P_s\} \subset \mathrm{R}[X_0,\ldots,X_k]$, of s homogeneous quadratic polynomials.
Output $b_k(S),\ldots,b_{k-\ell}(S)$, where S is the set defined by

$$S = \bigcap_{P\in\mathcal{P}} \{x \in \mathbf{S}^k \ \mid \ P(x) \leq 0\}.$$

Procedure
Step 1 Using Algorithm 2 compute the truncated complex $\mathcal{D}_\ell^{\bullet,\bullet}$, i.e.

$$\begin{array}{rcll} \mathcal{D}_\ell^{p,q} & = & \mathcal{D}^{p,q}, & 0 \leq p \leq \ell+1,\ \ k-\ell-1 \leq q \leq k, \\ & = & 0, & \text{otherwise}, \end{array}$$

Step 2 Compute using linear algebra, the ranks of

$$H^i(\mathrm{Tot}^\bullet(\mathcal{D}_\ell^{\bullet,\bullet})),\ \ k-\ell+1 \leq i \leq k.$$

Step 3 For each i, $k-\ell \leq i \leq k$, output, $b_i(S) = \mathrm{rank}(H^i(\mathrm{Tot}^\bullet(\mathcal{D}_\ell^{\bullet,\bullet})))$.

Complexity Analysis: The number of algebraic operations is clearly bounded by $\sum_{i=0}^{\ell+2}\binom{s}{i}k^{2^{O(\min(\ell,s))}}$ using the complexity analysis of Algorithm 2. □

Proof of Correctness: The correctness of the algorithm is a consequence of the correctness of Algorithm 2 and Theorem 6.2. □

Remark 7.1. Suppose that (using Notation from Algorithm 3) $\mathcal{P}' \subset \mathcal{P}$ and

$$S' = \bigcap_{P\in\mathcal{P}'} \{x \in \mathbf{S}^k \ \mid \ P(x) \leq 0\},$$

and letting $\mathcal{D}_\ell'^{\bullet,\bullet}$ denote the corresponding complex for S', it is clear from the definition that there is a homomorphism, $\Phi_{\mathcal{P},\mathcal{P}'} : \mathcal{D}_\ell^{\bullet,\bullet} \to \mathcal{D}_\ell'^{\bullet,\bullet}$ defined as follows.

For

$$\phi = \bigoplus_{\mathcal{Q}\subset\mathcal{P},\#\mathcal{Q}=p+1} \phi_{\mathcal{Q}} \in \mathcal{D}_\ell^{p,q} = \bigoplus_{\mathcal{Q}\subset\mathcal{P},\#\mathcal{Q}=p+1} \check{\mathcal{M}}_{\mathcal{Q}}^q,$$

$$\Phi_{\mathcal{P},\mathcal{P}'}(\phi) = \bigoplus_{\mathcal{Q}\subset\mathcal{P}',\#\mathcal{Q}=p+1} \phi_{\mathcal{Q}}.$$

Recall from 6.2 that there exists,

$$\psi : \mathcal{D}_\ell^{\bullet,\bullet} \longrightarrow \mathcal{N}_\ell^{\bullet,\bullet}$$

which induces an isomorphism, $\psi^* : H^*(\mathrm{Tot}^\bullet(\mathcal{D}_\ell^{\bullet,\bullet})) \longrightarrow H^*(\mathrm{Tot}^\bullet(\mathcal{N}_\ell^{\bullet,\bullet}))$.

Denoting by $\mathcal{N}'^{\bullet,\bullet}_\ell$ the (truncated) Mayer-Vietoris complex for S' and by $i_{\mathcal{P},\mathcal{P}'} : \mathcal{N}^{\bullet,\bullet}_\ell \to \mathcal{N}'^{\bullet,\bullet}_\ell$ the inclusion homomorphism, we have the following commutative diagram.

$$\begin{array}{ccc} H^*(\mathrm{Tot}^\bullet(\mathcal{D}^{\bullet,\bullet}_\ell)) & \xrightarrow{\Phi^*_{\mathcal{P},\mathcal{P}'}} & H^*(\mathrm{Tot}^\bullet(\mathcal{D}'^{\bullet,\bullet}_\ell)) \\ \downarrow{\scriptstyle \psi^*} & & \downarrow{\scriptstyle \psi'^*} \\ H^*(\mathrm{Tot}^\bullet(\mathcal{N}^{\bullet,\bullet}_\ell))) & \xrightarrow{i_*} & H^*(\mathrm{Tot}^\bullet(\mathcal{N}'^{\bullet,\bullet}_\ell))) \end{array}$$

Note that $H^*(\mathrm{Tot}^\bullet(\mathcal{N}^{\bullet,\bullet})) \cong H_*(S)$ and $H^*(\mathrm{Tot}^\bullet(\mathcal{N}'^{\bullet,\bullet})) \cong H_*(S')$.

It is clear that Algorithm 3 can be easily modified to output the complex $\mathcal{D}^{\bullet,\bullet}_\ell$, by outputting the matrices corresponding to the vertical and horizontal homomorphisms in the chosen bases. Furthermore, given a subset $\mathcal{P}' \subset \mathcal{P}$, Algorithm 3 can be made to output both the complexes $\mathcal{D}^{\bullet,\bullet}_\ell$ and $\mathcal{D}'^{\bullet,\bullet}_\ell$ along with the matrices defining the homomorphism $\Phi_{\mathcal{P},\mathcal{P}'}$ with the same complexity bounds.

8. The General Case

Let $\mathcal{P} = \{P_1, \ldots, P_s\} \subset \mathrm{R}[X_1, \ldots, X_k]$ with $\deg(P_i) \leq 2, 1 \leq i \leq s$, and let $S \subset \mathrm{R}^k$ be the basic semi-algebraic set defined by $P_1 \leq 0, \ldots, P_s \leq 0$.

Let $\varepsilon > 0$ be an infinitesimal, and let $P_{s+1} = \varepsilon \sum_{j=1}^k X_j^2 - 1$. Let $\tilde{S} \subset \mathrm{R}\langle\varepsilon\rangle^k$ be the basic semi-algebraic set defined by, $P_1 \leq 0, \ldots, P_s \leq 0, P_{s+1} \leq 0$.

Proposition 10. The homology groups of S and $\tilde{S}$ are isomorphic.

Proof. This is a consequence of the conical structure at infinity of semi-algebraic sets (see for instance Proposition 5.50 in [9])y. □

Moreover, denoting by P_i^h the homogenization of P_i, and $\tilde{S}^h \subset \mathbf{S}^k$ the set defined by the system of quadratic inequalities,

$$P_1^h \leq 0, \ldots, P_s^h \leq 0, P_{s+1}^h \leq 0,$$

on the unit sphere in $\mathrm{R}\langle\varepsilon\rangle^{k+1}$ we have,

Proposition 11. For $0 \leq i \leq k$, $b_i(\tilde{S}) = \frac{1}{2} b_i(\tilde{S}^h)$.

Proof. First observe that $\tilde{S}$ is bounded, and $\tilde{S}^h$ is the projection from the origin of the set $\tilde{S} \subset \{1\} \times \mathrm{R}\langle\varepsilon\rangle^k$ onto the unit sphere in $\mathrm{R}\langle\varepsilon\rangle^{k+1}$. Since, $\tilde{S}$ is bounded, the projection does not intersect the equator and consists of two disconnected copies in the upper and lower hemispheres, and each copy is homeomorphic to $\tilde{S}$. □

Algorithm 4 (Computing the top ℓ Betti Numbers: the general case).

Input A family of polynomials $\{P_1, \ldots, P_s\} \subset \mathrm{R}[X_1 \ldots, X_k]$, with $\deg(P_i) \leq 2$.
Output $b_{k-1}(S), \ldots, b_{k-\ell}(S)$, where S is the set defined by

$$S = \bigcap_{P \in \mathcal{P}} \{x \in \mathrm{R}^k \mid P(x) \leq 0\}.$$

Procedure
Step 1 Replace the family $\mathcal{P}$ by the family, $\mathcal{P}^h = \{P_1^h, \ldots, P_s^h, P_{s+1}^h\}$.
Step 2 Using Algorithm 3 compute $b_k(\tilde{S}^h), \ldots, b_{k-\ell}(\tilde{S}^h)$.

Step 3 Output $b_k(S) = \frac{1}{2} b_k(\tilde{S}^h), \ldots, b_{k-\ell}(S) = \frac{1}{2} b_{k-\ell}(\tilde{S}^h)$.

Proof of Correctness: The correctness of Algorithm 4 is a consequence of Propositions 10 and 11 and the correctness of Algorithm 3. □

Complexity Analysis: The complexity of the algorithm is clearly

$$\sum_{i=0}^{\ell+2} \binom{s}{i} k^{2^{O(\min(\ell,s))}}$$

from the complexity analysis of Algorithm 3. □

As in the case of Algorithm 3(see Remark 7.1) Algorithm 4 can be modified to output a complex whose associated total complex is quasi-isomorphic to the truncated Mayer-Vietoris complex of S. Furthermore, given a subset $\mathcal{P}' \subset \mathcal{P}$ defining the set $S' \supset S$, Algorithm 4 can be made to output the corresponding complexes of both sets, S and S', as well as the homomorphism between them induced by inclusion. Since this fact is important in certain applications we record it as a separate theorem.

Theorem 8.1. *There exists an algorithm, which takes as input a family of polynomials $\{P_1, \ldots, P_s\} \subset \mathrm{R}[X_1 \ldots, X_k]$, with $\deg(P_i) \leq 2$, and a number $\ell \leq k$, and outputs a complex $\mathcal{D}_\ell^{\bullet,\bullet}$. The complex $\mathrm{Tot}^\bullet(\mathcal{D}_\ell^{\bullet,\bullet})$ is quasi-isomorphic to $\mathrm{C}_\bullet^\ell(S)$, the truncated singular chain complex of S, where*

$$S = \bigcap_{P \in \mathcal{P}} \{x \in \mathrm{R}^k \mid P(x) \leq 0\}.$$

Moreover, given a subset $\mathcal{P}' \subset \mathcal{P}$, with

$$S' = \bigcap_{P \in \mathcal{P}'} \{x \in \mathrm{R}^k \mid P(x) \leq 0\}.$$

the algorithm outputs both complexes $\mathcal{D}_\ell^{\bullet,\bullet}$ and $\mathcal{D}'^{\bullet,\bullet}_\ell$ (corresponding to the sets S and S' respectively) along with the matrices defining a homomorphism $\Phi_{\mathcal{P},\mathcal{P}'}$, such that $\Phi^*_{\mathcal{P},\mathcal{P}'} : H_*(\mathrm{Tot}^\bullet(\mathcal{D}_\ell^{\bullet,\bullet})) \cong H_*(S) \to H_*(S') \cong H^*(\mathrm{Tot}^\bullet(\mathcal{D}'^{\bullet,\bullet}_\ell))$ is the homormorphism induced by the inclusion $i : S \hookrightarrow S'$. The complexity of the algorithm is $\sum_{i=0}^{\ell+2} \binom{s}{i} k^{2^{O(\min(\ell,s))}}$.

9. Hardness

In this section we show that the problem of computing the first few Betti numbers of a semi-algebraic set defined by quadratic inequalities is #P-hard. Note that PSPACE-hardness of the problem of counting the number of connected components for general semi-algebraic sets were known before [13, 26] and the proofs of these results extend easily to the quadratic case.

Recall that in the classical Turing machine model, a function taking its values in $\mathbb{N}$ belongs to the class #P, if there exists a polynomial time non-deterministic Turing machine, whose number of accepting paths on any particular input is equal to the value of the function on that input [25]. A function f is #P-hard if every function in #P can be reduced in polynomial time to the computation of f. The problem of counting the number of satisfying assignments to a Boolean formula is a #P-hard problem [25].

The results of this section complements the main algorithmic result of this paper proved in the last section, that the problem of computing the top few Betti numbers of such a set is in P.

Theorem 9.1. *Given a family of polynomials* $\mathcal{P} = \{P_1, \ldots, P_s\} \subset \mathrm{R}^k$*, with* $\deg(P) \leq 2, P \in \mathcal{P}$*, as input, the problem of computing* $b_\ell(S)$*, where* $S \subset \mathrm{R}^k$ *is defined by* $P_1 \geq 0, \ldots, P_s \geq 0$ *and* $\ell = O(\log k)$*, is #P-hard.*

Proof. We first prove the case when $\ell = 0$ by a straightforward reduction from the Boolean satisfiability problem. Clearly, if $\{C_1, \ldots, C_m\}$ is an instance of the Boolean satisfiability problem with clauses $C_1, \ldots, C_m$ and variables $Z_1, \ldots, Z_n$, then the number of satisfying assignments is equal to the number of connected components of the set defined by,

$$X_1(X_1 - 1) \geq 0, X_1(1 - X_1) \geq 0, \ldots, X_n(X_n - 1) \geq 0, X_n(1 - X_n) \geq 0$$
$$\bar{C}_1 - 1 \geq 0, \ldots, \bar{C}_m - 1 \geq 0,$$

where $\bar{C}_i \in \mathrm{R}[X_1, \ldots, X_n]$ is the linear polynomial obtained from the clause C_i by syntactically substituting all the disjunctions by additions, and for each $j, 1 \leq j \leq n$, the literal Z_j by the variable X_j, and the literal $\neg Z_j$ by the expression $(1 - X_j)$.

For $\ell > 0$, we reduce to the case $\ell = 0$ using the following observation. Given a basic semi-algebraic set $S \subset \mathrm{R}^k$ defined by, $P_1 \geq 0, \ldots, P_s \geq 0, \deg(P_i) \leq 2, 1 \leq i \leq s$, one can define another basic semi-algebraic set $S' \subset \mathrm{R}^N$ defined by M polynomials inequalities with degrees ≤ 2, with $M = s+4m+1$ and $N = k+2m+1$, where m is the total number of monomials in the polynomials $P_1, \ldots, P_s$. The number of monomials in the new system is bounded by $10m$. The semi-algebraic set S' has the homotopy type of the suspension ΣS. Moreover, the description of S' can be computed in polynomial time from the description of S. It follows from the basic properties of the suspensions (see [28]) that $b_1(\Sigma S) = b_0(S)$, which proves that computing $b_1(S)$ is also #P-hard. Iterating the construction, that is taking suspensions of suspensions ℓ times, and noting that $b_\ell(\Sigma^\ell(S)) = b_0(S)$, gives the result for $\ell = O(\log k)$ since the number of inequalities and variables increases only polynomially in s and k for $\ell = O(\log k)$.

We now describe the construction of the set S'. Introduce one new variable, X_0 and consider the semi-algebraic set, $S'' \subset \mathrm{R}^{k+1}$ defined by,

$$(1 - X_0^2)P_1 \geq 0, \ldots, (1 - X_0^2)P_s \geq 0,$$
$$1 - X_0^2 \geq 0.$$

The set $S'' = ([-1, 1] \times S) \cup H_1 \cup H_2$ where H_1 and H_2 are the hyperplanes defined by $X_0 = 1$ and $X_0 = -1$ respectively. It is easy to see that S'' has the same homotopy type as the suspension of S. However, the polynomials used in the description of S'' can have degrees as large as 4. We show below that by introducing new variables any quartic polynomial inequality can be reduced to a set of quadratic inequalities. The set S' is defined as follows. Write each monomial m of degree > 2 appearing in a polynomial used in the definition of S'', as a product of two monomials, m_1, m_2, each of degree at most 2. Now replace each ocurrence of m in the inequalities used in the definition of S'', by the quadratic monomial $Y_{m,1}Y_{m,2}$, where $Y_{m,1}, Y_{m,2}$ are new variables. Also, add the inequalities

$$Y_{m,1} - m_1 \geq 0, m_1 - Y_{m,1} \geq 0,$$
$$Y_{m,2} - m_2 \geq 0, m_2 - Y_{m,2} \geq 0,$$

to the set of inequalities. Clearly, the number of monomials, the number of new variables and the number of additional inequalities in the new system is bounded by $10m, 2m+1$ and $4m+1$ respectively, where m is the number of monomials in the original system. Finally, it is clear that the set S' defined by the above inequalities is a basic closed semi-algebraic set defined by polynomials inequalities of degree at most 2, which is homeomorphic to S'', and hence has the same homotopy type as ΣS. This completes the proof. □

References

[1] A.A. Agrachev, *Topology of quadratic maps and Hessians of smooth maps*, Algebra, Topology, Geometry, Vol 26 (Russian),85-124, 162, Itogi Nauki i Tekhniki, Akad. Nauk SSSR, Vsesoyuz. Inst. Nauchn.i Tekhn. Inform., Moscow, 1988. Translated in J. Soviet Mathematics. 49 (1990), no. 3, 990-1013.

[2] A.I. Barvinok, *Feasability Testing for Systems of Real Quadratic Equations*, Discrete and Computational Geometry, 10:1-13 (1993).

[3] S. Basu, *Computing the first few Betti numbers of semi-algebraic sets in single exponential time,* Journal of Symbolic Computation, Volume 41, Issue 10, October 2006, 1125-1154.

[4] S. Basu, *On Bounding the Betti Numbers and Computing the Euler Characteristics of Semi-algebraic Sets*, Discrete and Computational Geometry, 22:1-18 (1999).

[5] S. Basu, *On different bounds on different Betti numbers*, Discrete and Computational Geometry, 30:65-85 (2003).

[6] S. Basu, R. Pollack, M.-F. Roy, *Computing the first Betti number and the connected components of semi-algebraic sets*, preprint. Available at [arXiv:math.AG/0603248].

[7] S. Basu, R. Pollack, M.-F. Roy, *On the Combinatorial and Algebraic Complexity of Quantifier Elimination*, Journal of the ACM, 43, 1002-1045, (1996).

[8] S. Basu, R. Pollack, M.-F. Roy, *Computing Roadmaps of Semi-algebraic Sets on a Variety*, Journal of the AMS, vol 3, 1 55-82 (1999).

[9] S. Basu, R. Pollack, M.-F. Roy, Algorithms in Real Algebraic Geometry, Second Edition, Springer-Verlag, 2006.

[10] L. Blum, F. Cucker, M. Shub, S. Smale Complexity and Real Computation, Springer-Verlag, 1997.

[11] J. Bochnak, M. Coste, M.-F. Roy, Géométrie algébrique réelle, Springer-Verlag (1987). Real algebraic geometry, Springer-Verlag (1998).

[12] G. E. Bredon, Sheaf Theory, Springer-Verlag (1996).

[13] P. Burgisser, F. Cucker, *Counting Complexity Classes for Numeric Computations II: Algebraic and Semi-algebraic Sets*, Journal of Complexity, 22(2):147-191 (2006).

[14] J. Canny, *Computing road maps in general semi-algebraic sets*, The Computer Journal, 36: 504–514, (1993).

[15] A. Gabrielov, N. Vorobjov *Betti Numbers for Quantifier-free Formulae*, Discrete and Computational Geometry, 33:395-401, 2005.

[16] L. Gournay, J. J. Risler, *Construction of roadmaps of semi-algebraic sets*, Appl. Algebra Eng. Commun. Comput. 4, No.4, 239-252 (1993).

[17] D. Grigor'ev, D.V. Pasechnik, *Polynomial time computing over quadratic maps I. Sampling in real algberaic sets*, Computational Complexity, 14:20-52 (2005).

[18] D. Grigor'ev, N. Vorobjov, *Counting connected components of a semi-algebraic set in subexponential time*, Computational Complexity 2: 133-186 (1992).

[19] R. M. Hardt, *Semi-algebraic Local Triviality in Semi-algebraic Mappings*, Am. J. Math. 102, 291-302 (1980).

[20] J. Heintz, M.-F. Roy, P. Solernò, *Description of the Connected Components of a Semialgebraic Set in Single Exponential Time*, Discrete and Computational Geometry, 11, 121-140 (1994).

[21] J. McCleary A User's Guide to Spectral Sequences, Second Edition Cambridge Studies in Advanced Mathematics, 2001.

[22] J. Milnor, *On the Betti numbers of real varieties*, Proc. AMS 15, 275-280 (1964).

[23] O. A. Oleĭnik, *Estimates of the Betti numbers of real algebraic hypersurfaces*, Mat. Sb. (N.S.), 28 (70): 635–640 (Russian) (1951).

[24] O. A. Oleĭnik, I. B. Petrovskii, *On the topology of real algebraic surfaces*, Izv. Akad. Nauk SSSR 13, 389-402 (1949).
[25] C.H. Papadimitriou, Computational Complexity, Addison-Wesley (1994).
[26] J. Reif *Complexity of the mover's problem and generalizations*, Proceedings of IEEE Symposium on Foundations of Computer Science, 421-427, 1979.
[27] J. Renegar *On the computational complexity and geometry of the first order theory of the reals*, Journal of Symbolic Computation, 13: 255–352 (1992).
[28] E. H. Spanier Algebraic Topology, McGraw-Hill Book Company, 1966.
[29] R. Thom, *Sur l'homologie des variétés algébriques réelles*, Differential and Combinatorial Topology, 255–265. Princeton University Press, Princeton (1965).

School of Mathematics, Georgia Institute of Technology, Atlanta, GA 30332, U.S.A.
E-mail address: saugata@math.gatech.edu